ROBotic Open-architecture Technology for Cognition, Understanding and Behavior

Project no. 004370

RobotCub

Development of a cognitive humanoid cub

Instrument: Integrated Project
Thematic Priority: IST – Cognitive Systems

D5N.1 Imitation and Communication for the iCub

Due Date: Month 65
Submission date: Month 65

Start date of project: **01/09/2004** Duration: **65 months**

Organisation name of lead contractor for this deliverable: **EPFL**

Responsible Person: **Aude Billard**

Revision: **rev.no. 1.0**

Project co-funded by the European Commission within the Sixth Framework Programme (2002-2006)		
Dissemination Level		
PU	Public	**PU**
PP	Restricted to other programme participants (including the Commission Service)	
RE	Restricted to a group specified by the consortium (including the Commission Service)	
CO	Confidential, only for members of the consortium (including the Commission Service)	

D5N.1 Imitation and communication for the iCub

1.1 Introduction

This deliverable is the final report on work conducted within WP5 and WP6, on imitation and communication respectively. To recall, during the fifth year of Robotcub, these two workpackages were merged into a single workpackage named WP5N.

Works done within WP5N took complementary avenues to the understanding of imitation and communication in humans and robots. First, work done by UNIFE and EPFL, separately and in collaboration, analyzed different aspects of imitation and communication in humans (solely in humans). Second, work done by EPFL focused on the design of imitation learning in robots with an emphasis on improving robots' performance while paying no attention to how effective the method is from the human's view point. In contrast, work at the UNIHER looked at how robot control affects human-robot interaction and how the robot controller can be designed to improve human-robot communication and imitation.

The rest of this report is divided into the above main trends of research, namely study of imitation and communication in humans and design and study of controllers for human-robot imitation and communication. All of the work on which we report below has been published. Hence, only short abstract summarizing the goals and results of these studies are given with a complete reference to the associated publications. All associated publications are provided as pdf files in the attachments. All the work that relates to design of robot controllers for the iCub is available on the SVN repository.

1.2 Study of human imitation and communication

UNIFE and EPFL contributions to WP5N focused on the study of the pre-requisites to imitative and communicative behaviour in humans. We summarize each project next starting with contributions to a better understanding of the selective process during observation and imitation and moving to contributions related to the neural and behavioural basis of action observation.

1.2.1 Which articulations are controlled during imitation of meaningless and meaningful gestures?

UNIFE and EPFL have conducted jointly an experiment to verify if the position of articulations influences the way in which the same action is imitated (**Gesierich B., Canto R, Fabbri Destro M., Fadiga L., Finos L., Hersch M., Oliynyk A., Craighero L., *Study of kinematics and eye movements during imitation***). The study contrasted imitation of transitive and intransitive reaching movements with normal and unusual elbow elevation. It was hypothesis that imitation would proceed according to a hierarchy of goals in which reaching to the target in transitive motion would take precedence over replicating the particular orientaton of the elbow articulation during reaching, Analysis of the data confirmed this hypothesis in that all subjects reproduced correctly the motion of the end-effector, while solely a few of the subject (deemed "good imitators") reproduced both the unusual elbow elevation and the end-effector motion irrespective of whether they were imitating transitive versus intransitive motion. Measurement of eye-movements during observation of the demonstrated movements showed a correlation between scanning of the elbow motion during demonstration and being a good imitator. Following from this, a computational model that accounts for reaching movement in which both the end-effector and the elbow motion can be controlled concurrently was developed **(Just, A., Petreska, B., Billard, A., Craighero, L., D'Ausilio, A., Oliynyk, A. and Fadiga, L., Point-to-Point Unconstrained Gestures: Modeling Wrist and Elbow Trajectories, *Human Movement Science* (submitted), 2009)**

The model forms a particular instantiation of the dynamical systems controller developed at EPFL explaining the curvature of reaching motion **(Petreska and Billard, Biological Cybernetics, 2009)** and the generic framework for estimating non-linear dynamical systems motion models for robot control (**Hersch et al, IEEE Trans. In Robotics 2008, Gribovskaya et al Intern. Journal of Robotics Research , 2009 (submitted)**).

1.2.2 Functional magnetic resonance technique was used to verify brain activity during the exchange of gazes, which constitutes the simplest form of interpersonal communication.

UNIFE conducted a study of sympathy, considered as the ability of the observer to reproduce the internal states of others, either when observing an external event or the display of a reaction, motor or affective. The experimental question dealt with the implication that a subject has very little to know on his own internal states, so brain activity related to sympathy should be smaller than it is when a different subject is involved. Five different conditions have been used. The key comparisons were between the brain activity of a subject when he is looking at a different person and when he is looking at his own eyes. In other conditions, subjects were looking at an observer who was not looking, or they were looking at as they are not looked. A group of 29 subjects has been observed. Results support the hypothesis of sympathy as an information acquisition. For example, BA 44 (Broca's area) is involved specifically when two subjects exchange gazes but not when the subject is looking at himself. Anterior Insula is activated when subjects are being looked at and are not looking.

Data were presented in:
Rustichini, A.; Fadiga, L.; Lungu, O. Eye-to-eye communication. 469.7 2005 Neuroscience Meeting Planner. Washington, DC: Society for Neuroscience, 2005. Online.

1.2.3 An experiment has been performed demonstrating that Broca's area has a specific role not only in speech perception and production but also in others' action recognition.

Broca's area has been considered, for over a century, as the brain centre responsible for speech production. Modern neuroimaging and neuropsychological evidence have suggested a wider functional role is played by this area. In addition to the evidence that it is involved in syntactical analysis, mathematical calculation and music processing, it has recently been shown that Broca's area may play some role in language comprehension and, more generally, in understanding actions of other individuals. As shown by functional magnetic resonance imaging, Broca's area is one of the cortical areas activated by hand/mouth action observation and it has been proposed that it may form a crucial node of a human mirror-neuron system. If, on the one hand, neuroimaging studies use a correlational approach which cannot offer a final proof for such claims, available neuropsychological data fail to offer a conclusive demonstration for two main reasons: (i) they use tasks taxing both language and action systems; and (ii) they rarely consider the possibility that Broca's aphasics may also be affected by some form of apraxia. We administered a novel action comprehension test—with almost no linguistic requirements—on selected frontal aphasic patients lacking apraxic symptoms. Patients, as well as matched controls, were shown short movies of human actions or of physical events. Their task consisted of ordering, in a temporal sequence, four pictures taken from each movie and randomly presented on the computer screen. Patient's performance showed a specific dissociation in their ability to re-order pictures of human actions (impaired) with respect to physical events (spared). Our study provides a demonstration that frontal aphasics, not affected by apraxia, are specifically impaired in their capability to correctly encode observed human actions.

Results have been published in: **Fazio P, Cantagallo A, Craighero L, D'Ausilio A, Roy AC, Pozzo T, Calzolari F, Granieri E, Fadiga L. Encoding of human action in Broca's area (2009) Brain. 132(Pt7):1980-8.**

1.2.4 A behavioral study of cooperation and competition during human interaction demonstrated that even when the outcome of the play indicates perfect cooperation, electrophysiological measurements may reveal differences in attitudes and beliefs that guide social interaction.

We studied the behavior of 12 pairs of (normal, right-handed) undergraduate students while they were involved in a simple coordination game requiring motor interaction. Three experimental conditions were defined according to whether a monetary prize was given to both or only one player, if the couple was successfully completing the required assignment. Electromyographic potentials (EMG) were recorded from the right first dorsal interosseus muscle, a muscle critically involved in the motor task. We also collected written answers from standard questionnaires from which we constructed individual measures based on organized group interaction, social involvement and altruism. These measures of 'Altruism' were collected to estimate individual pro-social or altruistic attitudes and behavior. Consistently with a simple behavioral model, by which EMG signals may reveal subjects' personal concern (utility) associated to the given task, experimental evidence shows that individual average EMG signal was increasing when the players where expecting a monetary reward. When we split the subject pool into two sub samples (according to the measures of Altruism obtained from the questionnaire), we found that monetary incentives explain the level of subjects' EMG signal only in the sub sample characterized by low SC or Altruism. These findings seem to support the possibility that an electrophysiological measure, such as EMG recording, could reveal the most profound attitudes and believes that guide social interaction.

Results have been **submitted to Social Neuroscience**: **Roberto Censolo, Laila Craighero, Giovanni Ponti, Leonzio Rizzo, Luciano Fadiga Prosocial attitude modulates muscle activity in a simple coordination game.**

1.2.5 A transcranial magnetic experiment (TMS) demonstrated that motor structures provide a specific functional contribution to the perception of speech sounds.

Listening to speech recruits a network of fronto-temporo-parietal cortical areas. Classical models consider anterior (motor) sites to be involved in speech production whereas posterior sites are considered to be involved in comprehension. This functional segregation is challenged by action perception theories suggesting that brain circuits for speech articulation and speech perception are functionally dependent. Although recent data show that speech listening elicits motor activities analogous to production, it's still debated whether motor circuits play a causal contribution to the perception of speech. We administered TMS to motor cortex controlling lips and tongue during the discrimination of lip- and tongue articulated phonemes. We found a neurofunctional double dissociation in speech sound discrimination, supporting the idea that motor structures provide a specific functional contribution to the perception of speech sounds. Moreover, findings show a fine-grained motor somatotopy for speech comprehension.

Results have been published in:
D'Ausilio A, Pulvermüller F, Salmas P, Bufalari I, Begliomini C, Fadiga L. (2009) The motor somatotopy of speech perception. Curr Biol.;19(5):381-5.

1.2.6 A transcranial magnetic (TMS) experiment demonstrated that during other's actions observation the muscle-specific facilitation of the observer's motor system reflects the degree of muscular force that is exerted in an observed action.

Two separate TMS-experiments are reported in which corticomotor excitability was measured in the hand area of the primary motor cortex (M1) while subjects observed the lifting of objects with different weights. The type of action 'grasping and lifting the object' was always identical but the grip force varied according to the object's weight. In accordance to previous findings, activity of M1 was shown to modulate in a muscle-specific way, such that only those parts of M1 that control the specific muscles used in the observed lifting action, become increasingly facilitated. Moreover, the muscle-specific facilitation pattern of M1 was shown to modulate in accordance to the force requirements of the observed actions, such that corticomotor excitability was considerably higher for observing heavy object lifting compared to light object lifting. Overall, these results indicate that observed object grasping, requiring different force levels, is mirrored onto the observer's motor system in a highly muscle-specific manner, as measured in M1. The measured force-dependent modulations of corticomotor activity in M1 are hypothesised to be functionally relevant for the observer's ability to infer the observed grip force and consequently the weight of the lifted object.

Results have been submitted in **European Journal of Neuroscience:**
Alaerts Kaat, Senot Patrice, Swinnen Stephan, Craighero Laila, Wenderoth Nicole, Fadiga Luciano. Force requirements of observed object lifting are encoded by the observer's motor system: A TMS-study.

1.2.7 Furthermore, an additional experiment has been performed to test if cognitive cues are able to interfere with motor system facilitation during observation of grip force.

Motor Evoked Potentials (MEP) elicited by TMS stimulation of the First Dorsal Interosseous (FDI) muscle representation were measured during the observation of reach-grasp-lift actions upon 6 different objects: 1) transparent empty bottle (VisLight); 2) transparent full bottle (VisHeavy); 3) opaque empty bottle (HidLight); 4) opaque full bottle (HidHeavy); 5) opaque full bottle labeled light (LabLight); 6) opaque full bottle labeled heavy (LabHeavy). Light objects were 50 and heavy were 500g. This difference translated into clear different pattern of muscle contraction and kinematics. TMS was applied when this difference was found to be maximal, in a 100ms window after the beginning of lifting. Condition 1 and 2 afforded full knowledge of weight difference and kinematics information, 3 and 4 only kinematics, 5 and 6 no kinematics but cognitive cues (labels).

We found significant difference in MEPs amplitude in the 1-2 and 3-4 comparisons but no difference in 5-6, 1-3 and 2-4. Results show that the motor cortex does scale for the amount of muscle activity present in the observed action by analysing movement kinematics. Moreover all subjects reported the presence of only 5 objects (they recognised only one opaque) arguing for an implicit processing carried out by the motor system.

Results have been presented at FENS Forum 2008:
Senot P., D'Ausilio A., Franca M., Caselli L., Craighero L. & Fadiga L. Implicit coding of observed kinematics: the case of lifting object with different weights. ***FENS Abstr., vol.4, 123.26, 2008***

1.3 Design and study of controllers for human-robot imitation and communication

In this section, we summarize research conducted at UNIHER and EPFL, starting with work at EPFL on learning non-linear dynamical systems from demonstration for robust robot control and moving to the various UNIHER contributions on the study of gesture communication and its role in human-robot communication. Specifically, work at UNIHER in the final project phase followed three key lines of investigation a) the study of gesture communication and imitation in user studies with child and adult participants, b) the development of a computational architecture for development and learning in human-humanoid interaction, and c) the in depth analysis of human-robot interaction in a variety of human-robot interaction scenarios in order to illuminate aspects of timing, social cues, motor interference and well as possible benefits of such interaction in robot assisted play for children.

1.3.1 EPFL Work on Estimating Non-Linear Dynamical Systems of Motion

To compensate for the scarsity of human demonstrations, methods for imitation learning must be able to infer a motion model outside contexts covered by the demonstrations. In this work, we revisited the statistical approach developed in (Hersch et al 2008).to determine explicitly the region of application of the inferred model. This offers the ability to generate asymptotically stable motions in an area of the workspace that exceeds that covered by the demonstration.

Motion imitation requires reproduction of a dynamical signature of a movement, i.e. a robot should be able to encode and reproduce a particular path together with a specific velocity and/or an acceleration profile. Such a motion encoding is advantageous in that i) it allows to generalize a motion to unseen context; ii) it provides fast on-line replanning of the motion in the face of spatio-temporal perturbations; iii) it may embed different types of dynamics, governed by different attractors. (**E. Gribovskaya, M. Khansari and A. Billard, Learning the Nonlinear Multivariate Dynamics of Motion of Robotic Manipulators,** ***Intern. Journal of Robotics Research (submitted)).***

Code for estimating the dynamics of motion was implemented on the iCub and will be provided on the SVN repository by month 65 together with the video of the demonstrator.

1.3.2 Investigation of the effect of physical presence on human-humanoid gesture interaction games.

For this task, a study with human participants interacting with various degrees of physical presence was carried out. Results are reported in the journal paper "**Effects of Embodiment and Gestures on Social Interaction in Drumming Games with a Humanoid Robot**" **(Kose-Bagci et al, 2009**. Here, we present results from an empirical study investigating the effect of embodiment and minimal gestures in an interactive drumming game consisting of an autonomous child-sized humanoid robot (KASPAR) playing with child participants. Each participant played three games with a humanoid robot that played a drum whilst simultaneously making (or not making) head gestures. The three games included the participant interacting with the real robot (physical embodiment condition), interacting with a hidden robot when only the sound of the robot is heard (disembodiment condition; note that the term 'disembodiment' is used in this paper specifically to refer to an experimental condition where a physical robot produces the sound cues, but is not visible to the participants), or interacting with a real-time image of the robot (virtual embodiment condition). We used a mixed design where repeated measures were used to evaluate embodiment effects and independent-groups measures were used to study the gestures effects. Data from the implementation of a human–robot interaction experiment with 66 children are presented, and statistically analysed in terms of participants' subjective experiences and drumming performance of the human–robot pair. The subjective experiences showed significant differences for the different embodiment conditions when gestures were used in terms of enjoyment of the game, and perceived intelligence and appearance of the robot. The drumming performance also differed significantly within the embodiment conditions and the presence of gestures increased these differences significantly. The presence of a physical, embodied robot enabled more interaction, better drumming and turn-taking, as well as enjoyment of the interaction, especially when the robot used gestures. While the experiments in this study had to be conducted with KASPAR, the findings are also very relevant to future studies with the iCub in cases where different robot embodiments are used.

1.3.3 Investigation of using cues in the regulation of human-humanoid interaction games.

This direction formed a major part of UNIHER's scientific work. Resulting in eight publications in the final period of RobotCub, some of which represent the investigation into cues such as timing, gesture, body expression and communicative aspects of imitation (former Task 5.1).

Some of these publications report the culmination or updates of previous work (such as with timing and robot-human drumming interactions; investigation into the adaptive regulation of robot behaviour in response to human-robot interaction), while others overview the issues and developments over the course of recent years (e.g. in methodology of human-humanoid studies, or design and deployment of the low-cost minimally expressive humanoid robot KASPAR specifically used for human-robot interaction experiments in RobotCub prior to the availability of the iCub to the UNIHER team).

1.3.4 An Experimental Investigation of Interference Effects in Human-Humanoid Interaction Games

(Qiming Shen, Hatice Kose-Bagci, Joe Saunders, Kerstin Dautenhahn (2009) An Experimental Investigation of Interference Effects in Human-Humanoid Interaction Games. IEEE RO-MAN 2009, 18th IEEE International Symposium on Robot and Human Interactive Communication Sep. 27 - Oct. 2, 2009, Toyama International Conference Center, Japan.

[This also relates to the design, experimental study and analysis aspects of mirroring, timing, body expression and communicative aspects of imitation in child-robot interaction and using computational models of imitative interaction games (Task 5.1)]

Investigating how people respond to and relate to robots is a multifaceted scientific challenge. This paper reports on an experimental investigation concerning movement interference effects between a human and a robot. We compare results with that obtained by Oztop et al. [E. Oztop, D. W. Franklin, T. Chaminade, and G. Cheng (2005). "Human-humanoid interaction: is a humanoid robot perceived as a human?" in *International Journal of Humanoid Robotics* 2(4): 537-559], however, in our study we used a small child-sized robot (KASPAR) with an overall human-like appearance. The experiment was conducted with both child and adult participants who interacted with a small humanoid robot using arm waving behaviours. The experimental setup was designed to be less constrained than in [Oztop et al, 2005] with an emphasis on playful interaction. The experimental results did not show evidence for interference effects. This might be due to a more game-like and less constrained experimental environment or to the specific features of the robot or both. In addition to measurements of the variance of the movements, we investigated a measure for behavioural synchrony between human and robot movements based on the concept of information distance. The results of information distance analysis indicated that most of the human participants were affected by the robot's behavioural rhythms. While our experiments did not show a movement interference effect, we found behavioural adaptation of participants' movement timing to the robot's movements. Thus, the measure of behavioural synchrony that we introduced appears useful for complementing other measures (such as variance) previously used in the literature. This work is relevant for the field of gesture communication for humanoid robotics in general concerning questions of motor interference and motor coordination in human-humanoid interactions.

The next three papers relate to the investigation of mechanisms to adjust levels of play in real time as a response to styles of interaction of a robot with people (continuing and completing the former Task 6.5). This research investigates gesture communication and social cues in human-robot interaction in a robot-assisted play context, utilizing concepts developed in the RobotCub project in a therapeutic application domain:

1.3.5 Towards socially adaptive robots: A novel method for real time recognition of human-robot interaction styles.

Dorothée François, Kerstin Dautenhahn and Daniel Polani (2009) Using Real-Time Recognition of Human-Robot Interaction Styles for Creating Adaptive Robot Behaviour in Robot-Assisted Play. IEEE Symposium on Artificial Life 2009, ALIFE'09, Nashville USA, pp.45 – 52. ***[Winner of Best Paper Award]***

Automatically detecting different styles of play in human-robot interaction is a key challenge towards adaptive robots, i.e. robots that are able to regulate the interactions and adapt to different interaction styles of the robot users. In this paper we present a novel algorithm for pattern recognition in human-robot interaction, the Cascaded Information Bottleneck Method. We apply it to real-time autonomous recognition of human-robot interaction styles. This method uses an information theoretic approach and enables to progressively extract relevant information from time series. It relies on a cascade of bottlenecks, the bottlenecks being trained one after the other according to the existing Agglomerative Information Bottleneck Algorithm. We show that a structure for the bottleneck states along the cascade emerges and we introduce a measure to extrapolate unseen data. We apply this method to real-time recognition of Human-Robot Interaction Styles by a robot in a detailed case study. The algorithm has been implemented for real interactions between humans and a real robot. We demonstrate that the algorithm, which is designed to operate real time, is capable of classifying interaction styles, with a good accuracy and a very acceptable delay. Our future work will evaluate this method in scenarios on robot-assisted therapy for children with autism.

1.3.6 Using Real-Time Recognition of Human-Robot Interaction Styles for Creating Adaptive Robot Behaviour in Robot-Assisted Play

Dorothée François, Daniel Polani, Kerstin Dautenhahn (2009) Towards Socially Adaptive Robots: A Novel Method for Real Time Recognition of Human-Robot Interaction Styles. ***Proc 8th IEEE-RAS Int Conf on Humanoid Robots (Humanoids 2008), December 1-3*****, 2008, pp.353-359**

This paper presents an application of the Cascaded Information Bottleneck Method for real-time recognition of Human-Robot Interaction styles in robot-assisted play. This method, that we have developed, is implemented here for an adaptive robot that can recognize and adapt to children's play styles in real time. The robot rewards well-balanced interaction styles and encourages children to engage in the interaction. The potential impact of such an adaptive robot in robot-assisted play for children with autism is evaluated through a study conducted with seven children with autism in a school. A statistical analysis of the results shows the positive impact of such an adaptive robot on the children's play styles and on their engagement in the interaction with the robot.

A long-term study of children with autism playing with a robotic pet: Taking inspirations from non-directive play therapy to encourage children's proactivity and initiative-taking

Dorothée François, Stuart Powell and Kerstin Dautenhahn (2009). "A long-term study of children with autism playing with a robotic pet: Taking inspirations from non-directive

play therapy to encourage children's proactivity and initiative-taking". In: Robots in the Wild: Exploring human-robot interaction in naturalistic environments: Special Issue of *Interaction Studies* 10:3, pp. 324–373

This paper presents a novel methodological approach of how to design, conduct and analyse robot-assisted play. This approach is inspired by nondirective play therapy. The experimenter participates in the experiments, but the child remains the main leader for play. Besides, beyond inspiration from non-directive play therapy, this approach enables the experimenter to regulate the interaction under specific conditions in order to guide the child or ask her questions about reasoning or affect related to the robot. This approach has been tested in a long-term study with six children with autism in a school setting. An autonomous robot with zoomorphic, dog-like appearance was used in the studies. The children's progress was analyzed according to three dimensions, namely, Play, Reasoning and Affect. Results from the case-study evaluations have shown the capability of the method to meet each child's needs and abilities. Children who mainly played solitarily progressively experienced basic imitation games with the experimenter. Children who proactively played socially progressively experienced higher levels of play and constructed more reasoning related to the robot. They also expressed some interest in the robot, including, on occasion, affect.

Also relating to Task 5.1, the work using Drum-mate to study human-humanoid gestural and timing cues is reported in the three papers, highlighting the important role of interaction dynamics and timing in human-humanoid gesture communication and interaction:

Hatice Kose-Bagci, Frank Broz, Qiming Shen, Kerstin Dautenhahn, Chrystopher L.Nehaniv (2010), As Time Goes by: Representing and Reasoning Timing in the Human-Robot Interaction Studies, AAAI - Spring Symposium 2010: It's All in the Timing: Representing and Reasoning About Time in Interactive Behavior, Stanford University, Palo Alto, California. Extended Abstract (accepted for publication).

In **(Kose-Bagci et al, 2010),** we summarize the experimental design issues related to timing in three human-robot interaction studies investigating imitation and drumming experiences with child-sized humanoid robots and human participants. Here the focus is to understand the role of timing in interaction design. Rather than humanoid robotics that merely mimics human behaviours (e.g. waving or drumming), the goal is instead engage with human interaction partners in a 'social manner', e.g. in a call and response turn-taking interaction for timing plays a crucial role in the emergent interaction. In particular, this work is aimed at understanding the temporal cues in interaction as part of generally applicable methods for developmental robotics with emphasis on imitation and gesture communication.

1.3.7 Drum-mate: Interaction dynamics and gestures in human-humanoid drumming experiments

Hatice Kose-Bagci, Kerstin Dautenhahn, Dag S. Syrdal and Chrystopher L. Nehaniv (in press) Drum-mate: Interaction dynamics and gestures in human-humanoid drumming experiments. *Connection Science.*

This journal article **(Kose-Bagci et al, in press)** investigates the role of interaction kinesics in human–robot interaction (HRI). We adopted a bottom-up, synthetic approach towards interactive competencies in robots using simple, minimal computational models underlying the robot's interaction dynamics. We present two empirical, exploratory studies investigating a drumming experience with a humanoid robot (KASPAR) and a human. In the first experiment, the turn-taking behaviour of the humanoid is deterministic and the non-verbal gestures of the robot accompany its drumming to assess the impact of non-verbal gestures on the interaction. The second experiment studies a computational framework that facilitates emergent turn-taking dynamics, whereby the particular dynamics of turn-taking emerge from the social interaction between the human and the humanoid. The results from the HRI experiments are presented and analysed qualitatively (in terms of the participants' subjective experiences) and quantitatively (concerning the drumming performance of the human–robot pair). The results point out a trade-off between the subjective evaluation of the drumming experience from the perspective of the participants and the objective evaluation of the drumming performance. A certain number of gestures was preferred as a motivational factor in the interaction. The participants preferred the models underlying the robot's turn-taking which enable the robot and human to interact more and provide turn-taking closer to 'natural' human–human conversations, despite differences in objective measures of drumming behaviour. The results are consistent with the temporal behaviour matching hypothesis previously proposed in the literature which concerns the effect that the participants adapt their own interaction dynamics to the robot's.

1.3.7.1 Drumming with a Humanoid Robot: Lessons Learnt from Designing and Analysing Human-Robot Interaction Studies

In work relating also to tasks 5.1, 5.2, 6.1, 6.2, 6.3, 6.4, and 6.7 of the previous implementation workplan (now subsumed under Task 5N.2), we have summarized methodological and experimental design issues related to three human-robot interaction studies investigating a drumming experience with KASPAR, a humanoid child-sized robot, and (in total 116) human participants. Our aim is not to have KASPAR just replicate the human's drumming but to engage in a 'social manner' in a call and response turn-taking interaction. This requires the set up of enjoyable as well as (as much as possible) controlled experiments. Two Human-Robot Interaction (HRI) experiments with adult participants and one experiment with primary school children were carried out to investigate different aspects of such interactions. We briefly summarize issues concerning experimental methodology and design, as well as ethical, legal, safety issues in addition to many 'practical' challenges of setting up and conducting HRI experiments with an autonomous humanoid robot. This is reported in **Hatice Kose-Bagci, Kerstin Dautenhahn, and Chrystopher L. Nehaniv (2009), Drumming with a Humanoid Robot: Lessons Learnt from Designing and Analysing Human-Robot Interaction Studies, *Proc. AAAI - Spring Symposium 2009: Experimental Design for Real-World System*s, Stanford University, Palo Alto, California, March 22-25, 2009, AAAI Technical Report SS-09-03, AAAI Press, pp. 25-32.**

1.3.8 KASPAR – A Minimally Expressive Humanoid Robot for Human-Robot Interaction Research

Kerstin Dautenhahn, Chrystopher L. Nehaniv, Michael L. Walters, Ben Robins, Hatice Kose-Bagci, N. Assif Mirza, Michael Blow (in press) KASPAR - A Minimally Expressive Humanoid Robot for Human-Robot Interaction Research. Special Journal Issue on "Humanoid Robots" for journal *Applied Bionics and Biomechanic*s, published by Taylor and Francis.

This journal article **(Dautenhahn et al, in press)** provides a comprehensive introduction to the design of the minimally expressive robot KASPAR, developed and elaborated in the course of the RobotCub project and subsequently independently supported by the University of Hertfordshire. This was first undertaken as a design exercise to inform the iCub's design, including expressive features, and human-robot interaction aspects.

Since 2005, i.e. years before the iCub became available at UNIHER (only in 2009 as part of the FP7 IP ITALK), versions of KASPAR provided a humanoid platform in which to conduct research on imitation and gesture communication. The humanoid design is particularly suitable for human-robot interaction studies. A low-cost design with off-the-shelf components has been used in a novel design inspired from a multi-disciplinary viewpoint, including comics design and Japanese Noh theatre. The design rationale of the robot and its technical features are described in detail. The robot is particularly suitable for human-robot studies that are conducted in schools. Three research studies are presented in the article that have been using KASPAR extensively. Firstly, we present its application in robot-assisted play and therapy for children with autism. Secondly, we illustrate its use in human-robot interaction studies investigating the role of interaction kinesics and gestures. Lastly, we describe a study in the field of developmental robotics into computational architectures based on interaction histories for robot ontogeny. The three areas differ in the way how the robot is being operated and its role in social interaction scenarios. Each is introduced briefly and examples of the results are presented. Reflections on the specific design features of KASPAR that were important in these studies and lessons learnt from these studies concerning the design of humanoid robots for social interaction are discussed. An assessment of the robot in terms of utility of the design for human-robot interaction experiments concludes the paper.

1.3.9 Implementation of gesture communication interaction games integrated with the Interaction History Architecture (IHA) architecture demonstrable on a humanoid robot, release in the iCub software repository and demo on the iCub.

The IHA has been one of UNIHER's key contributions to the iCub software repository as a generic architecture for development and interaction histories developed for the iCub.

Frank Broz, Hatice Kose-Bagci, Chrystopher L. Nehaniv, Kerstin Dautenhahn, (2009), "Learning behavior for a social interaction game with a childlike humanoid robot", Social Learning in Interactive Scenarios Workshop, Humanoids 2009, Paris, France, 7 December, 2009.

In **(Broz et al. 2009)** we describe the integration of further multimodal social cues in IHA and its application for the iCub to acquire new sequences of actions in the context of social interaction and reinforcement by social cues such as visual gaze and involvement in turn-taking. In particular behaviours of peek-a-boo, building on our previous work (Mirza et al, 2007, 2008a; 2008b)[1], and

[1] N. A. Mirza, C. L. Nehaniv, K. Dautenhahn, and R. te Boekhorst (2007), "Grounded sensorimotor interaction histories in an information theoretic metric space for robot ontogeny," *Adaptive Behavior*, vol. 15, no. 2, pp. 167–187, 2007.
N. A. Mirza, C. L. Nehaniv, K. Dautenhahn, R. te Boekhorst (2008a), "Developing Social Action Capabilities in a Humanoid Robot using an Interaction History Architecture", *Humanoids 2008*, IEEE Press, 2008.
N. A. Mirza, C. L. Nehaniv, K. Dautenhahn, R. te Boekhorst (2008b), "Anticipating Future Experience using Grounded Sensorimotor Informational Relationships", *Artificial Life XI*, MIT Press, 2008.

drumming interactions (cf. the Kose-Bagci lead-author papers in the references, plus those from previous years of the RobotCub project).

Based on its sensorimotor experiences (as defined in previously delivered implementations of IHA and Mirza et al papers, plus deliverables D6.3 and D6.4), augmented by short term memory capacity, the humanoid is able to develop behaviours or other action sequences via social interaction and reinforcement via social cues. The state-of-the-art implementation potentially supports the behavioural development, switching between acquired behaviours, including acquisition of peek-a-boo and engagement in drumming interactions. This implements episodic memory and prospective action-selection within a constrained temporal horizon, providing functional support for several aspects of Cognitive Development (WP2) and allowing the robot potentially to scaffold its development of behavioural competencies, while making use of forgetting and merging of experiences.

Software development of this extended Interaction History Architecture, augmented by the use of other modules (e.g. updated AudioAnalyser, Drum-mate, gaze tracking) in the RobotCub software repository are demoed as part of D5N.2, and this software has been delivered as a software release in the iCub repository as part of Deliverable D5N.3 (which comprises updates and documentation of D5.7, D6.3, and other relevant modules).

1.3.10 Interaction History Architecture (IHA): Development in a Social Context (D5N.2 Demo)

UNIHER's final demo on the iCub is intended to show development of behaviour in an open-ended face-to-face social interaction game between the iCub and a human. This work is based on the Interaction History Architecture (IHA), and is an extension of the software that supported the earlier IHA peek-a-boo demo on the iCub to allow more types of interaction and social cues. The human participant is able to interact with the robot and provide it with positive social feedback using their presence and gaze direction, as well as by playing a drum. The robot uses this feedback to acquire behaviour that leads to sustained interaction with the human. The robot comes to associate sequences of simple actions and gestures (waving, hitting the drum, etc.) executed under certain conditions with successful interaction based on its
history of experience. Note that the robot selects from a set of predefined actions and behaviours. While these behaviours are high-level goal directed actions (e,g., "hit the drum"), the goals and effects of the action are not specified directly to the iCub. Instead, the robot learns which behaviours and sequences of behaviours are successful by evaluating their effects during interaction based on the reward they achieve, based on social drives. The extended Interaction History Architecture is intended to support the robot developing different socially communicative, scaffolded behaviours in the course of temporally extended social interactions with humans by making use of social drives and its own first-person experience of sensorimotor flow during social interaction dynamics.

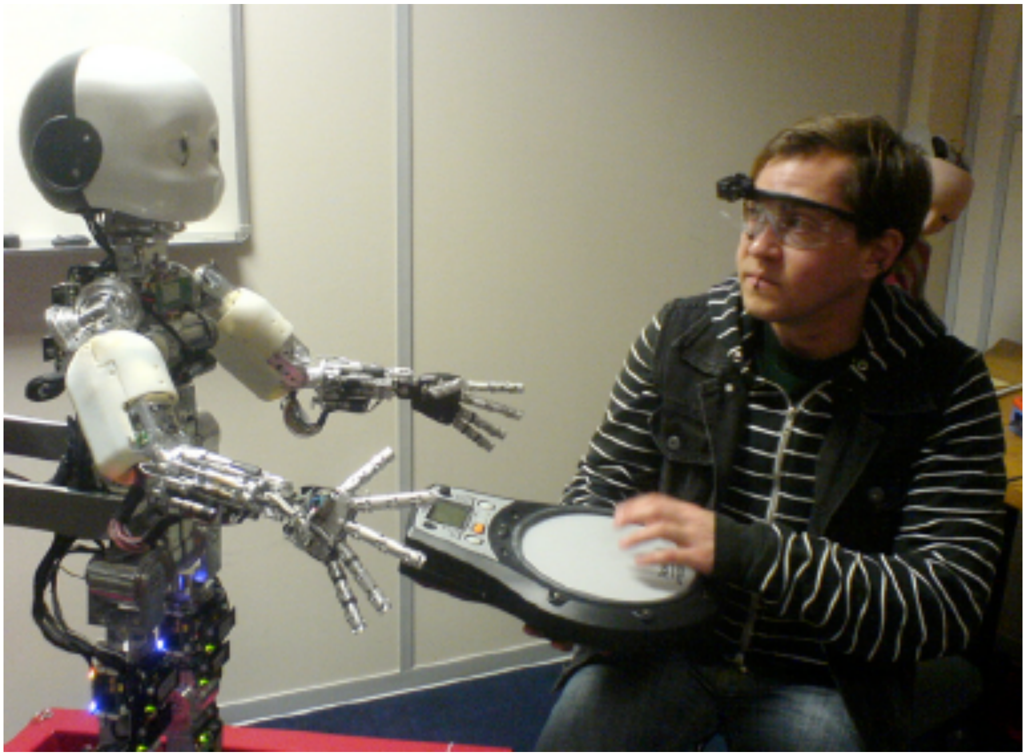

Current state of the art on-board gaze tracking in humanoid robotics does not yet allow for the robust, real-time, on-board detection of mutual gaze, gaze direction detection, or the framing of joint reference using deixis and gaze in relatively unconstrained human-robot interaction scenarios. Therefore the gaze-tracking (but not face-tracking or head-motion tracking) employed in conjunction with IHA at present involves use of a head-mounted tracker device for the human participant. This allows us to now already develop other gesture communication competencies that will be able to exploit such on-board technology, without waiting for the technology to become available. Studies using this type of set up will provide further detail and establish baselines for human-humanoid interaction. These capabilities for the iCub may soon become available on-board as the open-source repository grows and hardware continues to be upgraded by the iCub community, now much larger than the RobotCub Consortium. For example, the ITALK project focused on the emergence of linguistic communication based on integrating social and motor learning, and which has four iCubs, is collaborating with UNIHER to establish and exploit methods using gaze-tracking as an aspect of social gesture communication important in the acquisition of linguistic competencies by humanoid robots (including, e.g. new lexical items, holophrases, and grammatical constructions) via grounded sensorimotor interaction histories and social cues in interaction.

1.3.11 Updated Interaction History Architecture and Related Software Modules (contributing to software deliverable D5N.3).

In order to develop behaviours based on episodic experience in the sensorimotor domain and apply prospection, while scaffolding based on already mastered behaviours (Mirza et al 2007, 2008a, 2008b), and now also while integrating social cues for social learning and interaction, the Interaction History Architecture (IHA) has been substantially updated. The old IHA is still in the repository in the same place familiar to its existing user base. The new version of IHA and the applications that run it have been checked into the repository as a separate application, since their revisions make them incompatible with the older version that some users may still be relying on.

The new version and apps are in the repository under:
iCub/src/interactionHistoryNew iCub/app/ihaNew
Old version at: iCub/src/interactionHistory iCub/app/iha

There are several changes between the new and the old version of IHA motivated by issues of integration, software engineering optimization, as well as the modeling of development that takes into account multimodal social cues to scaffold ontogeny and learning. The structure of the architecture has been changed and streamlined for greater simplicity. In the old IHA (see the module diagram at http://eris.liralab.it/iCub/dox/html/group__icub__iha.html), the sensor data from different input modules went both into the sensorimotor processing module and the motivation dynamics module, then the output of the motivation dynamics module (the reward based on the sensor data) was passed into the sensorimotor processing module. In the new version, all sensing modules send their input only to the sensorimotor processing module. There, the sensor values are collected into a single data frame. This data is then passed to the motivation dynamics, where it is used to compute reward. Then the sensor values and their resulting reward are passed to the data storage module in a common frame. The code which determined the facial expression that the robot should display as feedback about its current reward state was moved from the sensorimotor processing module to the motivation dynamics module.

A diagram of the module architecture accompanies the on-line documentation. The IHA face detect module can now run in two modes, detect only (this is the version used in the old IHA) and tracking. The tracking mode uses openCV's Haar wavelet based face detector to find a face. Once one (or more) faces are detected, the largest found is used to initialize openCV's camshift algorithm, which does colour histogram-based tracking. This method is more robust to partial occlusions and changes of head orientation, as well as less computationally expensive than the face detection. Note that the face detector is used on two different input streams in the new application. In one, the human's face is detected in the robot's eye camera. The robustness of this is greatly improved by using the tracking mode. The other face detector runs on the scene camera from the gaze tracker and detects the iCub's own face in the human's field of view. The tracking mode is not used with this face detector, because the iCub's colour is not very distinct from most common backgrounds and the iCub's head does not significantly change orientation during the demo interaction (except when hiding during peek-a-boo, when it is also occluded).

Additionally, there were several different configuration files used by the application modules which listed redundant information (information about action definitions, for example) which could become out of sync and incorrect if modifications were made to one file and not the others. In the new IHA, such data is only present in one configuration file, which may be loaded by the multiple modules which use it.

There are several new modules added to the new version of IHA:

-mobileeye
This module compiles and displays the gaze tracker data. It takes as input: the scene image from the gaze tracker, the gaze direction (in scene image coordinates), and the output of a face detector run on the scene image. The current gaze direction and location of any faces are displayed on the scene image, and the face and gaze coordinates are sent to the sensorimotor processing module. The drivers which obtain the scene image and gaze coordinates from the gaze tracker we use

(ASL's Mobile Eye) contain proprietary source code, and are therefore not included in the repository. However, this module can be used with any gaze tracking system that reports similar output.

-shortTermMemory
This module aggregates certain streams of sensor data into a history window of user specified length. The data is sent from the sensorimotor processing module. The data history is used to create task-specific "scores" for each data channel that is watched. These scores are sent to the motivation dynamics module to be used in the computation of reward. The scores are computed based on the proportion of the time that the robot is observing certain events while executing turn-taking relevant actions (such as drumming or the hiding phase of peek-a-boo) over the duration of the history window.

-audioAnalyser
This is a modification of the audio analyser for Dr. Kose-Bagci's drummate code, modified to work with IHA. It connects to portaudio for sound input. The user can specify the rate at which messages about the occurrence and duration of drumbeats are sent.

Biological Cybernetics manuscript No.
(will be inserted by the editor)

A. Just · B. Petreska · A. Billard · L. Craighero · A. D'Ausilio · A. Oliynyk · L. Fadiga

Point-to-Point Unconstrained Gestures: Modeling Wrist and Elbow Trajectories

the date of receipt and acceptance should be inserted later

Abstract Although point-to-point reaching motions have received a lot of attention, the way these movements are controlled remains incompletely resolved. Different controllers seem to be recruited depending on the task. Unconstrained reaching movements in space are strongly curved, in opposition to the widely accepted view of quasi-straightness. We argue that the curvature of the movement is due to environmental constraints that affect directly the planning of the movement.

We propose a mathematical model whereby movements are planned through the combination of two concurrent controllers for the wrist and elbow in space. Coherence constraints are enforced between the two systems to simulate biomechanical constraints at the wrist, elbow and shoulder levels. External constraints, such as the presence of obstacles, are encapsulated in a virtual force which affects the planning of the movement.

The predictions of the model are validated against kinematic data from human reaching motions. Four types were contrasted: intransitive versus transitive reaching motions and natural versus un-natural motions. In the un-natural case, subjects were requested to exaggeratedly elevate the elbow during the movement. In all four movements types, the movements are highly curved. The model renders with high accuracy the kinematics of the movements and accounts for the curvature as an effect of the virtual force.

Keywords gesture modeling, VITE model, optimization with Lagrange, unconstrained and voluntary motion, multi-joint arm movement

A. Just · B. Petreska · A. Billard
Learning Algorithms and Systems Laboratory, Ecole Polytechnique Fédérale de Lausanne, 1015 Lausanne. Switzerland

L. Craighero · A. D'Ausilio · A. Oleynik · L. Fadiga
Section of Human Physiology, University of Ferrara, 44100 Ferrara. Italy

L. Fadiga
The Italian Institute of Technology, 16163 Genova. Italy.

1 Introduction

Much attention has been devoted to the study of point-to-point reaching movements, most of which focused on movements restricted to a plane. These studies highlighted several invariant features (Gibet et al 2004), such as quasi-straightness of the hand path from initial and target positions and the so-called bell-shaped velocity profile (Morasso 1981). Soon, such simple rules were questioned when considering unconstrained motions instead of the usual paradigm of constrained motions, or so-called compliant motions (Desmurget et al 1997). Indeed, the majority of the studies of point-to-point movements were highly constrained and required subjects to hold a hand-held cursor. Unconstrained motions, in contrast, refer to free motions of the hand. Results from unconstrained studies show that the spatio-temporal characteristics of compliant and unconstrained movements are fundamentally different. (Desmurget et al 1997) showed that movement duration was higher in the compliant condition than for unconstrained movements. Furthermore, path curvature was significantly higher for unconstrained motions. Hence, compliant and unconstrained motions involve different control strategies. Evidence supports the hypothesis that unconstrained motions are not following a straight line but are slightly curved. This hypothesis is further supported by (Boessenkool et al 1998) who states that trajectory curvature is an inherent property of unconstrained arm movements.

Another largely unresolved issue of motor control relates to the redundancy of the arm joints. A simple way to illustrate this is to consider the various postures that the arm can adopt to touch the same target. Several mathematical models have tried to answer this delicate question. Choosing between describing the kinematics of the arm in Cartesian coordinates or in joint angle space is a thorny problem and evidence comes in support of either of the two representations depending on the task (Flash and Hogan 1985; Rosenbaum et al 1995; Torres and Zipser 2002). To overcome this problem, the movements are often described more abstractly in terms of a global measure. This measure encodes the cost of each movement and the optimal movement is the one that minimizes this cost function. Cost functions may be defined using either kinematics or dynamic information on the movement.

Cost functions based on kinematic information deal with geometrical and temporal information: position, velocity, acceleration, etc. In (Flash and Hogan 1985), the cost function is defined as the square of the magnitude of the jerk (rate of change of acceleration) integrated over the entire movement. The minimum jerk model generates smooth hand trajectories which are straight and follow a bell-shaped velocity profile.

Cost functions based on dynamic information depend on the forces acting on the hand and arm. The minimum torque change model (Uno et al 1989) proposes as measure of performance the square of the first derivative of the torque integrated over the entire movement. In (Uno et al 1989) the model was compared to the minimum jerk model for unconstrained horizontal movements between two targets located in the sagittal plane. It was shown that the minimum torque change model and minimum jerk model were both predicting straight hand paths. However, for trajectories starting with the arm stretched sideways, the two models gave very different predictions. The minimum jerk model still predicted a straight-line hand paths whereas the trajectories predicted by the minimum torque model were gently curved, and thus more similar to observed human motion.

Other methods have been proposed to model the arm trajectories. Harris and Wolpert proposed the minimum variance theory (Harris and Wolpert 1998). Their model is based on the physiological assumption that the control signal is corrupted by noise. In the presence of this noise, the shape of the hand trajectory is selected so as to minimize the variance of the final arm position. In (Ogihara and Yamazaki 1999), the authors take a very different approach. They modeled the nervous system as a recurrent neural network. Given a goal position, the modeled nervous system was able to generate muscular activation signals used to move the hand to the target position. An interesting feature of this model is its ability to model the position of the whole arm. Most of the models presented previously were dealing mainly with the hand trajectory. A method has been proposed in (Kang et al 2003) to model the arm with its 4 DOFs. The arm trajectory is decomposed into intermediate positions. The model solves the joint angles for these positions by minimizing the sum of absolute value of all joints' torque work in each sub-path (trajectory between two via-positions). Their model unfortunately showed poor results for the adduction/abduction angle of the shoulder. Following this same idea, Gu et al. proposed the equilibrium point based model (Gu and Ballard 2006). The human arm motion can be seen as a sequence of short motion segments. Movements are generated by gradually shifting from one segment position to the next.

The models we have reviewed in the previous paragraphs are mostly dealing with compliant gestures or are modeling solely the hand path. Few of those have been designed to predict the evolution of movement of the entire arm, from start to target. In the present paper, we propose a method for generating the position of the entire arm for point-to-point motions. Further, since the elbow and hand locations are known, the whole arm configuration is determined, we model the control of the arm trajectory with two concurrent dynamical systems driving the hand and elbow separately, but coupled through kinematics constraints. We extend the biologically plausible VITE model (Bullock and Grossberg 1988), that describes a dynamical system to generate straight point-to-point trajectories in the Cartesian space. The extended VITE model we propose accounts for the observed curvature of the movement. Note that an extension of the VITE model that generate curved writing movements has already been proposed (Bullock et al 1993). The extension consisted in running three coupled VITE controllers to control the x-, y- displacements and wrist rotation of the hand, respectively. The curvature was the result of initiating each model at different starting times. An important disadvantage of this approach to model point-to-point movement is that it required a series of multiple arbitrary targets for each curvature change, which is not the case with the EFF-VITE model.

In order to validate the model, we conduct motion studies, in which unconstrained reaching motions are generated. Most of the literature has focused on the study of reaching movements directed at a target (Atkeson and Hollerbach 1985; Desmurget et al 1997; Magescas and Prablanc 2006). To determine if the curvature of the movement results from generating transitive (i.e. directed to a target) versus intransitive movements, we contrast two conditions in which subjects either reach for an object or do a reaching motion directed to no particular location on a table. We hypothesized that in both conditions the trajectories would be curved and argue that this curvature is necessary and fulfills two main goals: to avoid uncomfortable arm postures (for example, it is more natural to extend the elbow to the right during the motion than keeping a purely straight trajectory) and to encapsulate environmental constraints such as the presence of the table.

Furthermore, in order to better understand how the central nervous system manages to decouple the control of the upper and lower arms, when forced to do so, we investigated the kinematics of motion in which the elbow was forced to follow a trajectory more elevated than that found during natural reaching movements. (Koshland et al 2000) showed that, reaching during movements, the wrist exhibited similar characteristics as the proximal joints, demonstrating a coupling among the joints. We thus expected the curvature of the trajectories of the wrist also to increase as an effect of the exaggerated elevation of the elbow.

In Section 2 we describe the dynamical systems driving the elbow and wrist motions and explain how coherence constraints between the wrist and elbow are enforced in the model. Section 2.2 describes the experimental set-up and procedure followed during the motion studies. A comparative analysis of the model's predictions and human data is done in Section 3, followed by a discussion of the model's biological plausibility.

2 Materials and Methods

2.1 Description of the model

Our proposed approach is based on an extension of Bullock and Grossberg's Vector Integration To Endpoint (VITE) model (Bullock and Grossberg 1988). The VITE model is a biologically inspired model that can only generate straight point-to-point trajectories. Contrary to the VITE model, the extended force-field version of the VITE model (EFF-VITE) can account for curved reaching movements, and can be used to model both the trajectories of the hand and elbow. Compared to the VITE model, the EFF-VITE model is time-independent and thus stable in case of long lasting perturbations. Furthermore, it represents a proper force governed system. In the EFF-VITE system, the trajectory of the hand or elbow is governed by the following dynamical system:

$$\ddot{x}(t) = \alpha(-\dot{x}(t) + \beta g(t)^{\delta}(h(t)+\gamma)(\frac{x^*(t)-x(t)}{\|x^*(t)-x(t)\|} + g(t)\mathbf{F}(t))) \quad (1)$$

and

$$\mathbf{F}(t) = g(t)\mathbf{u} + h(t)\mathbf{v} \quad (2)$$

where

$$g(t) = \frac{\|x^*(t)-x(t)\|}{\|x(t)-x(0)\| + \|x^*(t)-x(t)\|}$$

$$h(t) = \frac{\|x(t)-x(0)\|}{\|x(t)-x(0)\| + \|x^*(t)-x(t)\|}$$

are respectively the ratios between the distance separating the hand from the final target position x^* and the distance separating the hand from the initial position $x(0)$ over their total. The force $\mathbf{F}$ helps to comply with environmental constraints due to the volume and geometry of the body. $\mathbf{F}$ is the weighted sum of two constant force vectors that push the trajectory away from the straight line. $\mathbf{u}$ is the modulated force that perturbs the beginning of the movement, whereas $\mathbf{v}$ perturbs the end of the movement (Figure 1). The parameter $\alpha \in \mathbb{R}^+$ was fixed to a constant value. Parameters β, γ and δ control the general form of the velocity profile. β controls the asymmetry and peak value of the velocity profile. γ enables the initiation of the movement, and δ controls the final approaching phase of the movement and parameterizes the trade-off between precision and execution time. For example, lowering the value of δ shortens the movement deceleration phase but also increases the risk of overshooting the target position (Figure 2). The role of the parameters will be further discussed in Sections 3.2.2 and 3.2.3.

An arm configuration corresponds to a particular position in space of both the wrist and elbow. In the duo-EFF-VITE model, two concurrent EFF-VITE models are modeling the hand and elbow paths. As the hand and elbow are linked, these two systems are not independent. Hence, coherence constraints must be enforced in order to have a meaningful representation of the movement. Figure 3 presents the overall structure of the duo-EFF-VITE model. The outcome of the model is the position of the hand and elbow in the Cartesian space at each time step.

Let $\mathbf{x_w}$ and $\mathbf{x_e}$ be the position of the wrist and elbow in the 3D space where the origin is centered on the shoulder. The position of the arm is such that:

$$||\mathbf{x_e}|| = L_1 \quad (3)$$

and

$$||\mathbf{x_e} - \mathbf{x_w}|| = L_2 \quad (4)$$

where L_1 and L_2 are respectively the length of the upper-arm and forearm , and $||.||$ defines the vector norm.

Let $\mathbf{x_w}^d(t)$ and $\mathbf{x_e}^d(t)$ be the desired position of the wrist and elbow given by the EFF-VITE models at each time step t. In general, the variables $\mathbf{x_w}^d$ and $\mathbf{x_e}^d$ will not be consistent with kinematic constraints. In order to have consistent values, we find the values $\mathbf{x_w}^*$ and $\mathbf{x_e}^*$ that minimize the similarity measure H:

$$H(\mathbf{x_w}^*, \mathbf{x_e}^*) = ||\mathbf{x_w}^* - \mathbf{x_w}^d|| + ||\mathbf{x_e}^* - \mathbf{x_e}^d|| \quad (5)$$

under constraints given by equations (3) and (4).

The problem is solved analytically by using Lagrange optimization. We define the Lagrangian as:

$$L(\mathbf{x_w}^*, \mathbf{x_e}^*, \lambda_1, \lambda_2) = H + \lambda_1^T(||\mathbf{x_e}^*|| - L_1) + \lambda_2^T(||\mathbf{x_e}^* - \mathbf{x_w}^*|| - L_2) \quad (6)$$

To solve $\nabla L = 0$, we derive respectively $\frac{\partial L}{\partial \mathbf{x_w}^*}$, $\frac{\partial L}{\partial \mathbf{x_e}^*}$:

$$2(\mathbf{x_w}^* - \mathbf{x_w}^d) + \lambda_2||\mathbf{x_e}^* - \mathbf{x_w}^*||^{-1}(\mathbf{x_w}^* - \mathbf{x_e}^*) = 0 \quad (7)$$

$$2(\mathbf{x_e}^* - \mathbf{x_e}^d) + \lambda_1||\mathbf{x_e}^*||^{-1}\mathbf{x_e}^* + \lambda_2||\mathbf{x_e}^* - \mathbf{x_w}^*||^{-1}(\mathbf{x_e}^* - \mathbf{x_w}^*) = 0 \quad (8)$$

We thus need to solve the following system:

$$\begin{cases} 2(\mathbf{x_w}^* - \mathbf{x_w}^d) + \lambda_2||\mathbf{x_e}^* - \mathbf{x_w}^*||^{-1}(\mathbf{x_w}^* - \mathbf{x_e}^*) = 0 \\ 2(\mathbf{x_e}^* - \mathbf{x_e}^d) + \lambda_1||\mathbf{x_e}^*||^{-1}\mathbf{x_e}^* \\ \quad +\lambda_2||\mathbf{x_e}^* - \mathbf{x_w}^*||^{-1}(\mathbf{x_e}^* - \mathbf{x_w}^*) = 0 \\ ||\mathbf{x_e}^*|| - L_1 = 0 \\ ||\mathbf{x_e}^* - \mathbf{x_w}^*|| - L_2 = 0 \end{cases}$$

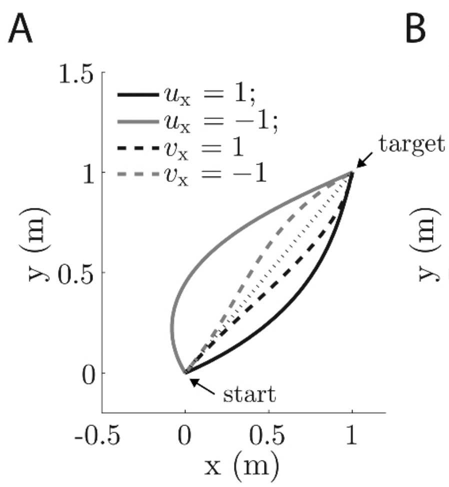

Fig. 1 Dynamics of the movement as a function of the force parameters. A: Forces are modulated such that u affects mostly the beginning of the movement and v mostly the end of the movement. The direction of the deviation from the straight trajectory is determined by the sign of the force. B: By combining the two forces u and v, trajectories that change direction can be obtained. Pararameter values: $\alpha = 50$, $\beta = 10$, $\gamma = 0.01$ and $\delta = 1$.

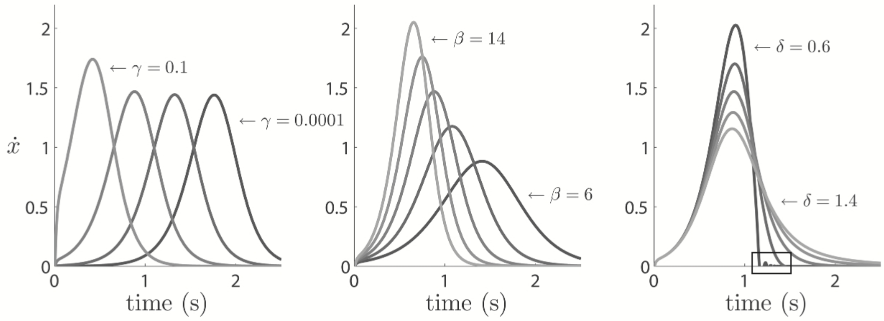

Fig. 2 Effect of the parameters γ, β and δ on the speed profile of the movements. The parameters γ (left) affects the beginning of the movement. The lower its value, the more time it takes the subject to start a movement. β (middle) controls the asymmetry and peak value of the velocity profile (α in our model is constant). δ (right) defined the approaching speed and thus parameterizes the trade-off between precision and execution time. In the rectangle, one can see the arm reaching the target too quickly and overshooting it at $\delta = 0.6$. Parameter values: $\alpha = 50$, $\beta = 10$, $\nu = 1.5$, $\gamma = 0.01$ and $\delta = 1$.

(9)

As the system has several solutions, we choose the solutions $\mathbf{x_w}^*, \mathbf{x_e}^* \in \mathbb{R}$ that minimize H. As the system is non-linear due to the presence of the norm, solutions are found numerically.

2.2 Experiments

Subjects Eight healthy subjects (4 females, 4 males, mean age 26 ±4) volunteered to perform a one-handed task consisting of point-to-point motions. All subjects were right-handed (Edinburgh Handedness Test, Oldfield (1971)) They were all naive regarding the purpose of the experiment. They reported no history of neurological or musculo-skeletal disorder. All had normal or corrected to normal vision.

Procedure Subjects sat comfortably on a chair in front of a table. They were asked to maintain a steady trunk position all along the recording session. Each hand movement started in the same rest position, with the forearm lying on the table and perpendicular to the trunk (Figure 4, left). Subjects were shown the movements by a demonstrator. There were two conditions. In the first condition, movements were directed towards an object placed 30 cm away from the subject in the sagittal plane (Figure 4, right). In the second condition, subjects had to reach in front of them and land their hand palm-down on the table. No location on the table was specified in this second condition. We refer to these two conditions

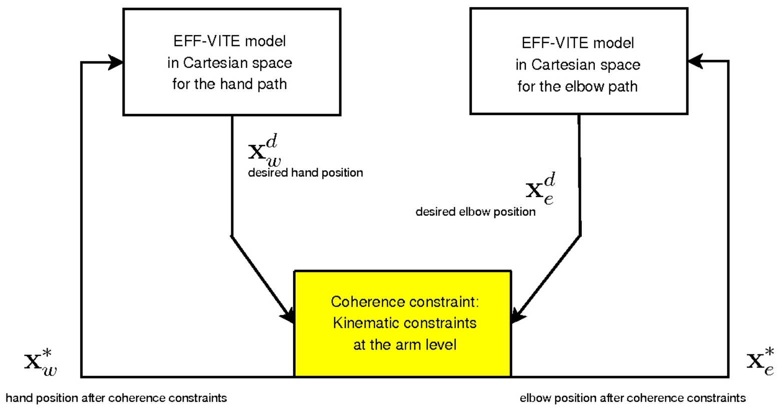

Fig. 3 The wrist-and-elbow path controller: The first EFF-VITE model (on the left) models the trajectory of the wrist in cartesian coordinates, whereas the second EFF-VITE model is used to model the elbow path in cartesian space. The coherence constraints ensure the desired positions $\mathbf{x}_w^d$ and $\mathbf{x}_e^d$ given by the EFF-VITE models are consistent relative to kinematic constraints. The modified values after coherence constraint for both the wrist and elbow positions, $\mathbf{x}_w^*$ and $\mathbf{x}_e^*$, are fed back to the EFF-VITE models.

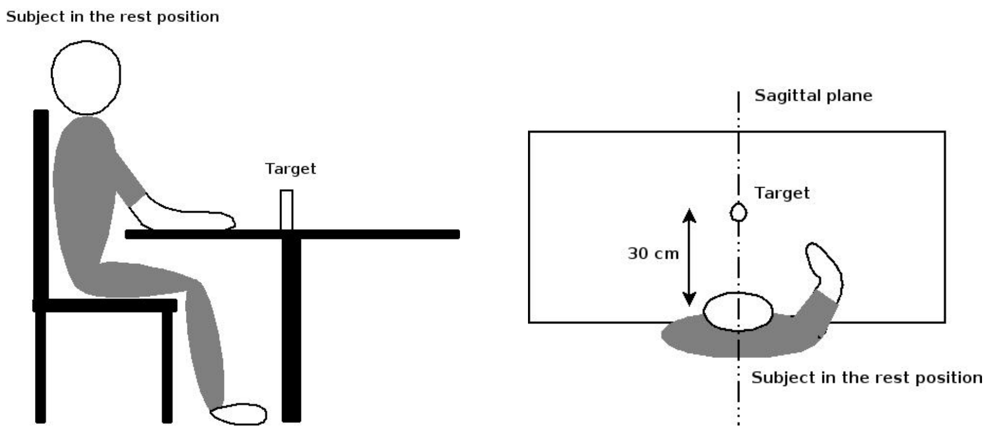

Fig. 4 Left: Experimental set-up seen from the right side with the subject in the rest position. Right: upper view of the set-up showing the position of the target when subjects performed transitive motions.

respectively as transitive (Trans) and intransitive (Intrans) movements in the rest of the paper.

For each condition, the subjects were instructed to perform two variants of the movements. In the first variant (so-called “Elb”), the subjects were asked to exaggeratedly elevate the elbow throughout the motion. In the second variant (so-called “Norm”), subjects were asked to perform motion in the way that seemed most natural to them. Movements were thus of four types: intransitive with normal kinematics (**Intrans Norm**), intransitive with an exaggerated elevation of the elbow (**Intrans Elb**), transitive with normal kinematics (**Trans Norm**) and transitive with an exaggerated elevation of the elbow (**Trans Elb**). Figure 5 presents snapshots of the four types of reaching movements.

Subjects were shown several times each movement types. Additional explanation was given when necessary. The subjects were instructed to replicate as precisely as possible these movements. A series of five movements for each condition and variant was recorded for each subject (Table 1).

Data acquisition The trajectory in space of the shoulder, elbow and wrist were recorded by using a kinematics recording system formed by three ProReflex MCU1000 cameras (QUALISYS AB, Sweden) detecting the 3D position of infrared reflecting markers (n=4) positioned on the left and right shoulders, right elbow and right wrist. The position of the markers was recorded at a frequency of 200 Hz during the execution of the movements. Figure 6 presents one subject wearing the markers as well as the shoulder-centered frame of reference used in the following of the paper to calculate wrist and and elbow trajectories.

Fig. 5 Snapshots of the four gesture types. From left to right: Intransitive action with normal kinematics and with an exaggerated elevation of the elbow. Transitive movement with normal kinematics and with an exaggerated elevation of the elbow. One can see that for the “Elb” variant the elbow position is always higher than for movements performed with normal kinematics for both the “Intrans” and “Trans” conditions.

Subjects	Repetitions	Recording sessions
8	5×4 gesture types	1

Table 1 Statistics of the database.

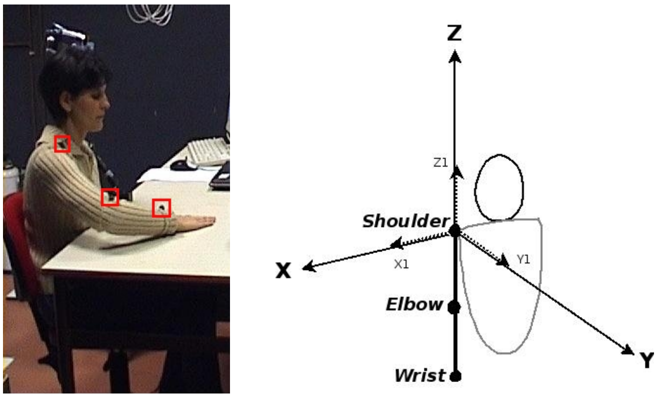

Fig. 6 Left: subject wearing markers on the right arm (markers are surrounded by red squares). Right: shoulder-centered frame of reference.

Data analysis All analyzes were performed using the Qualisys Track Manager (QUALISYS AB, Sweden) software, plus some custom programs written in Matlab (Mathworks, Natick, MA). Analysis was done solely on the reaching phase of each movement (from the rest position to the target location in the case of transitive movements, and from the rest position to the hand placement on the table in front of the subject for intransitive movements). Data were first segmented manually to remove any irrelevant movement prior to the onset of the reaching motion. We used only unfiltered raw values. The curvature index is computed as the ratio between the total arc length of the path and the Euclidian distance between the initial and final positions. A curvature index of 1 indicates a perfectly straight trajectory whereas a semi-circular path would have a curvature index of $CI = \pi/2$. The values of the model's parameters were optimized for each trial using 5^3 factorial experimental designs coupled with a local search procedure (Neter et al 1996; Hoos and Stützle 2004).

3 Results

3.1 Movement statistics

We first assessed the general characteristics of the recorded movements. For each movement type (Intrans Norm, Intrans Elb, Trans Norm, and Trans Elb), we computed the duration of the movement, path length and curvature index of the wrist and elbow on average across the 8 subjects and 20 trials (Table2).

Consistent with (Bernstein 1967)'s observations of substantial trial-to-trial variations, a three-way ANOVA analysis across subjects (eight levels), conditions (intransitive, transitive) and variants (elbow normal, elbow elevated) revealed a high *inter-subject* variability for both the duration of the movements, the length of the wrist path and the curvature index ($p < 0.001$), with a significant interaction effect for the subject/condition and subject/variant factors ($p < 0.01$ in each case, see Table 2). This high across subjects variability in performing the same motion is illustrated in Figures 7 and 8. Subject 9 tended to be very consistent across trials and

	Duration (s)	Path length (cm)		Curvature index		Elbow elevation (cm)
		Wrist	Elbow	Wrist	Elbow	
Intrans Norm	0.89 ± 0.28	25.3 ± 3.3	26.7 ± 3.5	1.16 ± 0.10	1.19 ± 0.06	-15.0 ± 2.3
Intrans Elb	1.11 ± 0.28	31.8 ± 5.7	37.2 ± 9.6	1.54 ± 0.28	1.52 ± 0.26	-7.0 ± 2.8
Trans Norm	0.84 ± 0.19	22.5 ± 3.0	23.0 ± 3.1	1.16 ± 0.09	1.16 ± 0.05	-15.9 ± 2.0
Trans Elb	1.14 ± 0.22	31.8 ± 6.2	34.9 ± 8.1	1.61 ± 0.45	1.47 ± 0.26	-6.4 ± 2.4
p-value (sub.)	< 0.001	< 0.001	< 0.001	< 0.001	< 0.001	< 0.001
p-value (cond.)	n.s.	< 0.003	< 0.001	n.s.	< 0.02	n.s.
p-value (var.)	< 0.001	< 0.001	< 0.001	< 0.001	< 0.001	< 0.001
p-value (sub*cond)	< 0.001	< 0.001	< 0.001	< 0.001	< 0.002	< 0.001
p-value (sub*var)	< 0.001	< 0.001	< 0.001	< 0.001	< 0.001	< 0.001
p-value (cond*var)	< 0.006	< 0.002	n.s.	n.s.	n.s.	< 0.001

Table 2 Duration, path length, curvature index and elbow elevation across trials and subjects. Three-way ANOVA showed that the movements performed with an exaggerated elevation of the elbow lasted longer, had a longer path for both the wrist and elbow and were significantly more curved than movements with normal kinematics. Furthermore, the recorded movements differed significantly accross subjects in their duration, path length, curvature index, and elbow elevation. The maximal height of the elbow during the movement was also significantly different accross the two motion variants.

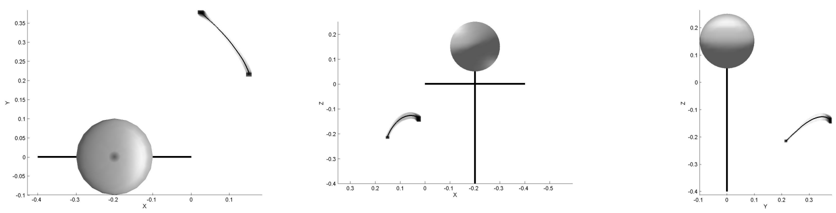

Fig. 7 Mean wrist trajectory (in black) and standard deviation envelope (in grey) for a transitive movement with an abnormal elevation of the elbow (Trans Norm) showing a small intra-variability for Subject 9.

displayed a low across trials variability of the wrist's motion (Figure 7), whereas Subject 5 displayed an overall much higher variability for the same motion (Figure 8). Given that the subjects had different arm lengths, the length of the wrist path varied importantly across subjects, especially in the intransitive case (see table 2).

All movements were curved ($CI > 1$). Most importantly for the argument of this paper, both the trajectory of the wrist and of the elbow were curved. The curvature is even more important for movements performed with an exaggerated elevation of the elbow ($CI > 1.6$). As a result, movements performed with an abnormal elevation of the elbow in both conditions (Intrans versus Trans) take significantly more time and are longer than movements performed with normal kinematics. Moreover, intransitive motions were significantly longer than transitive motions. This is likely due to the rotation of the wrist that occurs during intransitive motions (to place the palm down on the table), particularly when the movement is performed with an exaggerated elevation of the elbow (first two images in Figure 5).

3.2 Accuracy of the model

We measured the accuracy of the model to reproduce each instance of each motion type. We computed the mean deviation (MD) of the predicted wrist and elbow trajectories compared to the wrist/elbow trajectories at each time step, as well as the mean squared error (MSE) for each condition and variant of the movements. Table 3 provides these values for each gesture type. We also performed a three-way ANOVA analysis on these results for the subject, condition and variant factors. These results show no significant influence of either factor on the MSE for the wrist. For the elbow, the ANOVA analysis reveals a significant difference between the two motion variants (F=4.52, $p < 0.04$). However, the error is small and can be explained by the high variability of movements performed with an exaggerated elevation of the elbow (Elb variant).

Thus, overall, the model reproduces motions with high accuracy. It encapsulates the generic shape of both the trajectory in space and the speed profile of the wrist and elbow (Figure 9). 81% of the data for the wrist and 79% of the observed data for the elbow are reproduced by the model with a MSE inferior to the mean MSE. 3 to 4% of the errors are due to outlier data whereas another 53% are due to a poor reproduction of the start and/or end of the trajectory (Figure 10).

This is due to the fact that, like the original VITE model, the duo-EFF-VITE model, pre-supposes a smooth and gradually increasing and decreasing speed profile at the start and end of the movement, respectively. Because

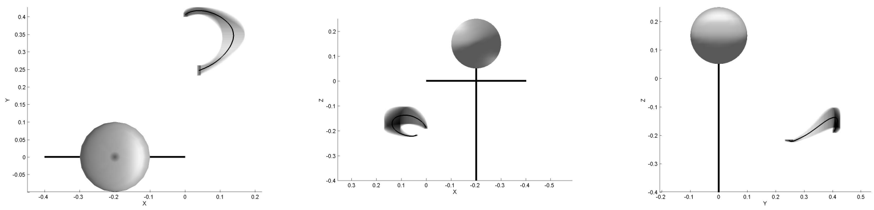

Fig. 8 Mean wrist trajectory (in black) and standard deviation envelope (in grey) for a transitive movement with an abnormal elevation of the elbow (Trans Norm) showing a high intra-variability for Subject 5.

	MD (cm)		MSE (cm^2)	
Movement	Wrist	Elbow	Wrist	Elbow
All motions	1.1 ± 0.7	1.1 ± 0.7	1.26 ± 5.16	1.09 ± 3.23
Intrans Norm	0.9 ± 0.5	0.9 ± 0.4	0.76 ± 1.32	0.65 ± 1.21
Intrans Elb	1.3 ± 0.4	1.3 ± 0.5	1.22 ± 0.86	1.15 ± 1.04
Trans Norm	0.8 ± 0.4	0.7 ± 0.4	0.48 ± 0.86	0.46 ± 1.05
Trans Elb	1.3 ± 1.1	1.4 ± 1.0	2.58 ± 10.10	2.09 ± 6.08
p-value (sub.)	< 0.02	< 0.002	n.s.	n.s.
p-value (cond.)	n.s.	n.s.	n.s.	n.s.
p-value (var.)	< 0.001	< 0.001	n.s.	< 0.04
p-value (sub*cond)	n.s.	n.s.	n.s.	n.s.
p-value (sub*var)	n.s.	n.s.	n.s.	n.s.
p-value (cond*var)	n.s.	n.s.	n.s.	n.s.

Table 3 Mean Deviation (MD) and Mean Squared Error (MSE) for the duo-EFF-VITE models on the trajectories of the wrist and elbow for each gesture type. We also provide three-way ANOVA results across subjects, movement conditions, variants, and interaction of these factors for each error type.

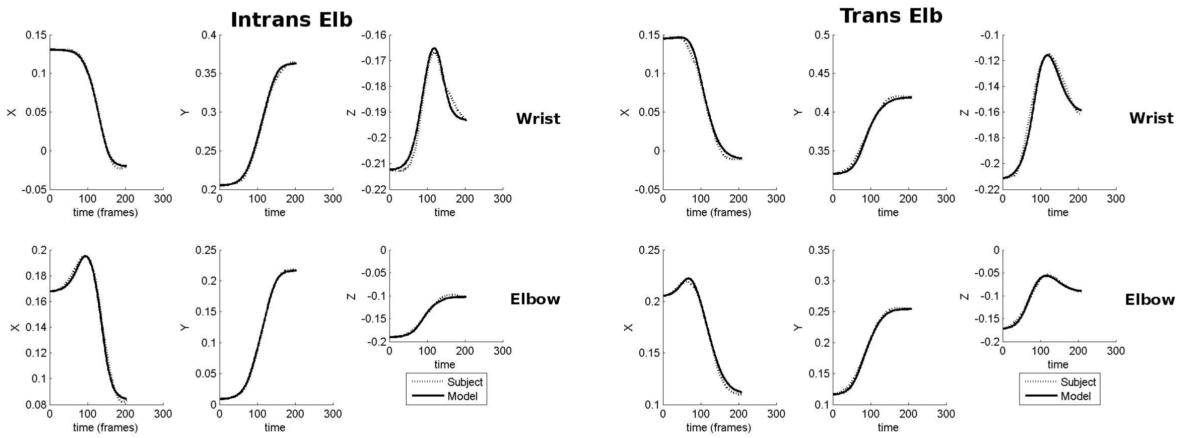

Fig. 9 Examples of movements well reproduced by the duo-EFF-VITE model.The trajectory of the subject's wrist (dotted line) and the modeled trajectory (black) are presented on top.

data were segmented manually, the speed profile was sometimes truncated and hence did not follow the typical pattern. Furthermore, some data present an atypical curvature at the start or end of the movement, due to hesitations on the subjects' parts. Because these imprecisions were minor and did not affect the generic characteristics of each motion (curvature and overall 3D spatial displacement), which we wanted the model to encapsulate, we did not eliminate the data.

3.2.1 Statistics of the model's parameters

A three-way ANOVA across subjects, conditions and variants, on the values taken by the force parameters of the model reveals that, while for the same subject the parameters for the wrist and elbow motions are consistent across conditions and variants, they vary importantly across subjects (see Tables 5 and 6). An effect of the variant (Norm versus Elb) is observed for the parameters

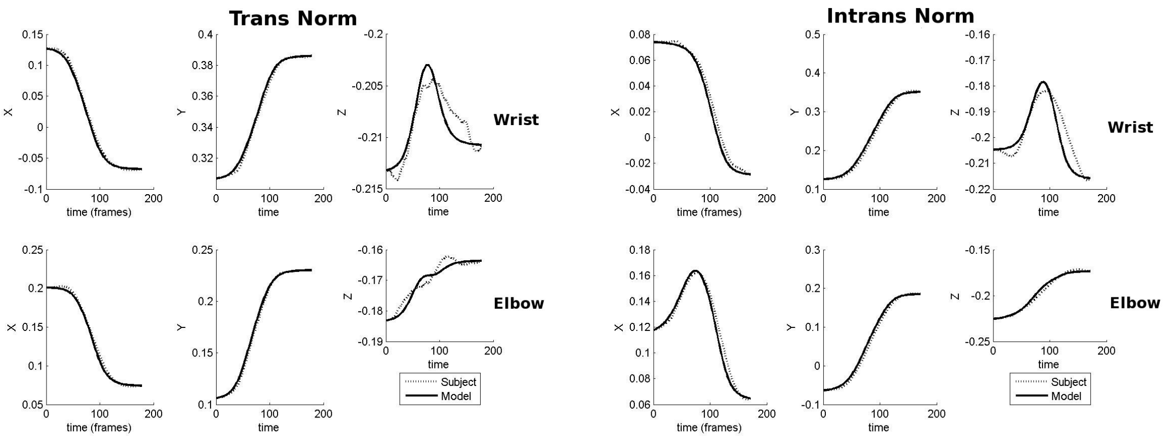

Fig. 10 Examples of movements poorly reproduced by the duo-EFF-VITE model. The trajectory of the subject's wrist (dotted line) and the modeled trajectory (black) are presented on top.

driving the elbow and this accounts for the variability with which subjects produced the required exaggerated elevation of the elbow (variability is expected given that the arm moved in an unconstrained manner).

We also computed the intra-subject variability of the wrist controller for movements with normal kinematics (Tables 8 and 9). We see that some subjects are more consistent in their movements than others, for both the force applied on the wrist and the parameters modulating the speed profile. This is particularly true for Subjects 6 and 8. This confirms the information contained in Figures 7 and 8, and is consistent with the general observation of a high inter-subject and inter-trial variability when performing the same motion, as discussed above and revealed in Table 2.

3.2.2 Meaning of the model's parameters

The parameters β, γ and δ in Equation 1 control the velocity profile of the movement. A two-way ANOVA shows that β and γ are similar across conditions and subjects (Table 4 in Annex) for the wrist controller. β controls the asymmetry and peak value of the velocity profile and γ determines the onset of the movements (Figure 2). As any irrelevant movement prior to the onset of the reaching motion has been manually removed, it is expected that γ takes a similar value across subjects and conditions. β is not significantly different across subjects, conditions and variants. Trajectories of the wrist thus follow the same velocity profile for both conditions (Intrans versus Trans) and variants (Norm versus Elb). δ controls the approaching speed of the movement. Together with β, δ determines a trade-off between overshooting the target and minimizing the execution time. Figure 11 presents the distribution of the values for β and δ for all movements. We see that the values are comprised within a region that minimizes execution time while ensuring a good precision of the movement.

3.2.3 Effect of the forces

We have already seen in Table 2 that the trajectories of both the wrist and elbow are curved. This curvature is accounted for by the values taken by the force parameters of the model (Tables 5 and 7). For each condition and variant of the movement, a non-null force is applied on the wrist and elbow. While one could have performed a straight-line motion in the normal condition, it is obvious that a straight path controller could not be envisioned for movements performed with an exaggerated elevation of the elbow. And, as expected, we observed larger values for the force parameters in the Elb variant of the movement.

The force applied along the x and y axes can also be related to the environmental and geometric constraints implied by the task. In our experiments, subjects sat on a chair with the body close to the table, the forearm resting on the table (Figure 6). To perform the movement, subjects needed to avoid the table ("table avoidance" constraint). To satisfy this constraint, the arm had to be placed above the table. Since the elbow is linked to the trunk by the upper-arm, all the possible positions of the elbow are located on a sphere centered on the shoulder and of radius the length of the upper-arm. Thus when the elbow tries to avoid the table, the elbow is also pulled away from the body along the x- and y-directions. Forces applied on the x- and y-axes are thus explained by the geometry of the body as well as the environmental constraints ("table avoidance").

The force along the z-axis (u_z and v_z) is close to zero in the "Norm" variant.However, in the "Elb" variant, the force along the z-axis at the end of the movement (v_z) (Table 7) is significantly higher (F=254.3, $p < 0.001$), with a mean value close to 1, so as to pull the elbow up during the motion. This effect is illustrated in Figures 13 and 12.

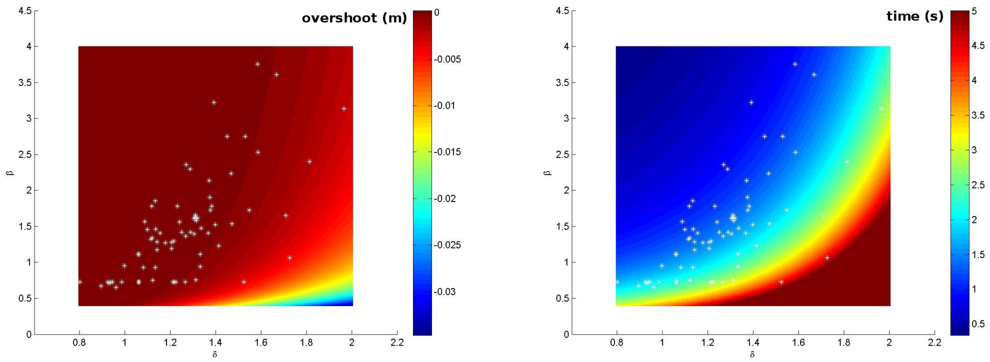

Fig. 11 Distribution of the parameters β and δ of the wrist controller, with respect to the overshoot distance (left) and execution time (right) for a 0.2 m movement. $\alpha = 50$, $\gamma = 0.02$

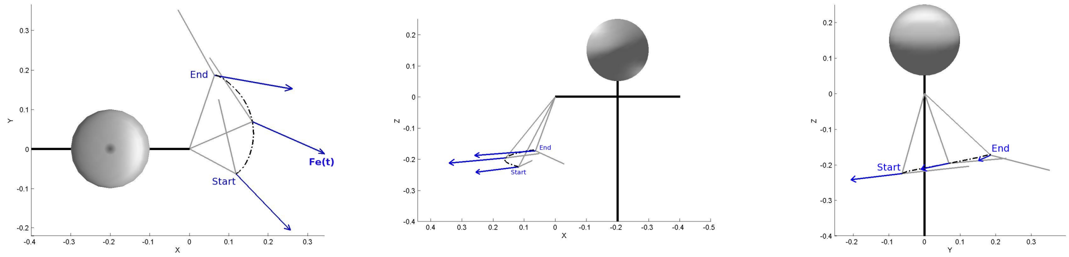

Fig. 12 Example of the force $\mathbf{F}_e$ applied on the elbow for an intransitive movement with normal kinematics. From left to right: projection in the xy-, xz- and yz-planes

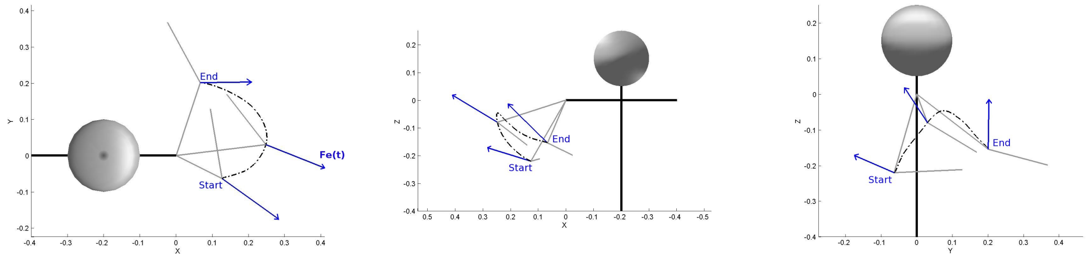

Fig. 13 Example of the force $\mathbf{F}_e$ applied on the elbow for an intransitive movement with an abnormal elevation of the elbow. From left to right: projection in the xy-, xz- and yz-planes

3.2.4 Separate controllers for wrist and elbow

As the elbow and wrist are linked by the forearm, the curvature of the hand path for movements performed with normal or exaggerated elevation of the elbow can be seen as a side effect of the elbow itself. Such correlation is revealed by looking at the Pearson coefficient between the forces [1] $\mathbf{F}_w$ and $\mathbf{F}_e$ (Equation (2)) applied on the wrist and elbow. These coefficients are respectively: $\rho(x) = 0.70$, $\rho(y) = 0.74$, and $\rho(z) = 0.18$, where $\rho(x)$, $\rho(y)$, and $\rho(z)$ are the Pearson coefficients along the x-, y- and z-axis, respectively. These results show that there exists a strong correlation between the force applied on the wrist and elbow along the x- and y-axis. The curvature of the wrist trajectories along the x- and y-axis is thus a side-effect of the elbow motion, and would contribute to confirm a view in which elbow and wrist are controlled by a single controller. In contrast, the wrist and elbow seem to be quasi-independent along the z-axis. This indicates that for the Elb variant of the movements, an exaggerated elevation of the elbow results in an increase in the amplitude of the virtual force $\mathbf{F}_e$ along the z-axis of the elbow controller only, and thus speaks in favor of

[1] The Pearson coefficient is the sum of the products of the normalized values of the two measures divided by the degree of freedom. The Pearson coefficient ranges from +1 to -1. If $\rho = 0$, then there is no linear relationship between the two variables. On the contrary, if $|\rho| = 1$, then there is a perfect linear relationship between the two variables.

having two separate controllers for the wrist and elbow, albeit correlated by geometrical constraints.

4 Discussion

4.1 Accuracy of the model

In this paper, we presented a model of reaching movements, which we validated against kinematic data of known motions in two conditions (intransitive versus transitive motions) and for two variants (movements performed with "naturally" versus movements performed with an exaggerated elevation of the elbow). We proposed an extension of the VITE model to account for both the curvature of naturally reaching movements and for the dual control of the wrist and elbow during unnatural reaching movements. The model gave an accurate account of the kinematics of the data for all the four movement types (Intrans Norm, Intrans Elb, Trans Norm and Trans Elb). Discrepancies between the model's prediction and the data for the velocity profiles at the start and end of the movement were observed in about 10% of the data. Closer analysis revealed that these errors were due to the fact that manual segmentation led to abrupt speed profiles, but also to the fact that in some cases, especially in transitive motions, the speed at the end of the reaching motion was not null (as subjects were transiting directly to a motion in which they grasped and lifted up the object). By construction, the duo-EFF-VITE model, like the VITE model, predicts a zero velocity at target. In effect, when transiting across two motions, subjects tend to displace the target of the reaching motion. One way to simulate this would be to introduce a new target position (corresponding to the final location of the subject's arm one the object had been lifted) slightly before the hand reached the original target point.

As expected, we observed significant inter-subjects and inter-trials variability across motions. To avoid these, we considered computing and modeling the mean trajectories of the wrist and elbow to capture the nature intrinsic to each movement independently from the subject and trial. This was ruled out as the mean movements of the wrist and elbow could no longer be correlated (since the correlations are not linear). Given that one of the hypotheses of the duo-EFF-VITE model is that the position of the wrist and elbow are controlled via two separate controllers acting in parallel but linked through biomechanical constraints, the effect of these biomechanical constraints would have been lost if we had worked with the mean trajectories. Besides, modeling each motion's instance allowed us to demonstrate that the curvature at the wrist level cannot be explained without taking into account the movement of the elbow.

4.2 Interpretation of the Model's Parameters

Parameters of the model are of two types. Three parameters β, γ, and δ are used to modulate the speed profile of the movement. They respectively control the general form of the velocity profile (asymmetry and peak value), enable the initiation of the movement and control the final approaching phase of the movement. Although the model's parameters were optimized to model each instance of the movements, we observed a consistency across the values of the parameters and showed that the parameter controlling the shape of the speed profile at the end of the movement takes values that optimize a trade-off between the precision and execution time of the whole movement. This is in agreement with the observation of a correlation across speed and accuracy of goal-directed movements (Plamondon and Alimi 1997; Meyer et al 1988). (Meyer et al 1988) hypothesized that this trade-off permits to cope optimally with noise in the human system.

Most importantly, the model hypothesized the existence of virtual forces that encapsulate tasks constraints to modulate a basic controller for reaching movements. We showed that these forces could explain the curvature of the movements of the wrist and elbow and could be interpreted in relation to environmental and biomechanical constraints. Further experiments should be conducted to validate this hypothesis by varying the task constraints, e.g. asking subjects to perform reaching motions by exaggeratedly lowering the elbow, and showing how the forces change as an effect of the context.

4.3 Separate Control of Wrist and Elbow

A second hypothesis inherent to the model is that elbow and wrist are driven by separate controllers, albeit correlated through imagined biomechanical constraints. Such a hypothesis corresponds to assuming that the nervous system is able to plan the mechanical effects that could arise from the motion of the arm segments (Galloway and Koshland 2001). An analysis of the relationship across the forces applied on the wrist and elbow at each time step revealed a strong correlation along the x- and y-axes. The forces along the z-axis were however quasi-independent of the elbow's elevation. The absence of correlation along the z-direction suggests that the motions of the wrist and elbow are computed separately by the brain. These conclusions are consistent with findings on multi-joint arm movements and with the Leading Joint Hypothesis (LJH) (Dounskaia et al 1998; Dounskaia 2005). The LJH states that there is one leading joint that guides the motion of the entire limb. Muscles of the secondary joints thus just play a regulatory role to ensure that the end-effector performs the required task. Interestingly, the LJH is applicable to our results if we consider the elbow as the leading joint and the wrist as the secondary joint.

4.4 Neural Correlated to the Model's Parameters

Similarly to the VITE model, the duo-EFF-VITE model depends on knowing at all time the wrist and elbow positions and velocities. Evidence that the velocity and position of the wrist may be explicitly computed and used for motor control by the nervous system exists. For instance, cells in the primary motor cortex (M1) of the monkey showed a high correlation between their discharge and the velocity profile of reaching movements (Moran and Schwartz 1999). Moreover (Wang et al 2006) confirmed the existence of a neural representation of the hand location in the motor cortex during reaching. They showed that position and velocity of the hand are simultaneously encoded by cortical motor neurons. Existence that the position and velocity of the elbow are explicitly computed is still questioned (Murphy et al 1982; Scott et al 1997; Reina et al 2001). While the duo-EFF-VITE model proposes a solution to encapsulate environmental and biomechanical constraints, it does not explain how the brain computes such constraints. As they contribute in several ways to the virtual forces, several brain areas may be involved.

Finally, the duo-EFF-VITE model is based on the idea that motions are not planned but unfold through time as the result of the inherent dynamics of the controllers. Such an approach is in line with the force-field approach (Graziano et al 2005), where the target of the motion acts as an attractor for the end-effector. Moreover, the model assumes that control is done in close-loop, taking into account the current position of the arm to correct the motion. This is supported by evidence that the nervous system is able to estimate and anticipate the state of the limb by integrating delayed sensory input and motor output, through afferent and efferent internal feedback loops (Desmurget et al 1997).

While the model exploits a representation of biomechanical constraints in the coupling of the elbow and wrist controllers, it does not account for the way the command are translated into muscle activation of the upper and lower arm limbs. While a complete understanding of the neural control of movements would require a realistic musculoskeletal model[2], we omitted such complexity in order to focus on explaining the gross dynamics of motor control. In particular, we aimed at explaining how volitional control of one specific limb (upper arm) could be done separately from that of the lower arm, as in the exaggerated elbow elevation condition considered here.

Movements presented in this paper were unconstrained. While this resulted in a high variability across trials and subjects' motions, it offered the opportunity to observe features of motion that are inherent to natural reaching motions. The duo-EFF-VITE model is however generic and could also model constrained movements. To confirm the LJH hypothesis and the use of the duo-EFF-VITE model in support of the latter, it would thus be interesting to replicate the present study with movements of the wrist constrained in the plane. The wrist would then become the leading joint and the elbow the follower. Results of such a comparative study would contribute to explaining the difference in the curvature of the hand path found for constrained and unconstrained movements (Desmurget et al 1997).

[2] Such model is very complex and difficult to obtain due to the numerous muscles and tendons present in the human arm (Cheng and Loeb 2008).

Acknowledgements This work was supported by the Sport and Rehabilitation Engineering Program at EPFL, CE Grants Robotcub, CONTACT and Poeticon, and Italian Ministry of University PRIN.

References

Atkeson C, Hollerbach J (1985) Kinematic features of unrestrained vertical arm movements. Journal of Neuroscience

Bernstein N (1967) The coordination and regulation of movements. Pergamon

Boessenkool J, Nijhof EJ, Erkelens C (1998) A comparison of curvatures of left and right hand movements in a simple pointing task. Experimental Brain Research

Bullock D, Grossberg S (1988) Neural dynamics of planned arm movements: emergent invariants and speed-accuracy properties during trajectory formation. Psychological Review

Bullock D, Grossberg S, Mannes C (1993) A neural network model for cursive script production. Biological Cybernetics

Cheng E, Loeb G (2008) On the use of musculoskeletal models to interpret motor control strategies from performance data. Journal of Neural Engineering

Desmurget M, Jordan M, Prablanc C, Jeannerod M (1997) Constrained and unconstrained movements involve different control strategies. The American Physiological Society

Dounskaia N (2005) The internal model and the leading joint hypothesis: implications for control of multi-joint movements. Experimental Brain Research

Dounskaia N, Swinnen S, Walter C, Spaepen A, Verschueren S (1998) Hierarchical control of different elbow-wrist coordination patterns. Experimental Brain Research

Flash T, Hogan N (1985) The coordination of arm movements: An experimentally confirmed mathematical model. The Journal of Neuroscience

Galloway J, Koshland G (2001) General coordination of shoulder, elbow and wrist dynamics during multijoint arm movements. Experimental Brain Research

Gibet S, kamp JF, Poirier F (2004) Gesture analysis: Invariant laws in movement. In: Gesture-based Communication in Human-Computer Interaction

Graziano M, Aflalo T, Cooke D (2005) Arm movements evoked by electrical stimulation in the motor cortex of monkey. Journal of Neurophysiology

Gu X, Ballard D (2006) An equilibrium point based model unifying movement control in humanoids. In: Robotics: Science and Systems

Harris C, Wolpert D (1998) Signal-dependent noise determines motor planning. Nature

Hoos H, Stützle T (2004) Stochastic local search: Foundations and applications. Morgan Kaufmann

Kang T, Tillery S, He J (2003) Determining natural arm configuration along reaching trajectory. In: Proc. of the Inter-

	β	γ	δ
Intrans Norm	2.04 ± 1.94	0.010 ± 0.005	1.31 ± 0.22
Intrans Elb	1.76 ± 1.53	0.006 ± 0.005	1.30 ± 0.37
Trans Norm	2.33 ± 1.75	0.011 ± 0.007	1.43 ± 0.30
Trans Elb	1.61 ± 1.36	0.007 ± 0.004	1.21 ± 0.39
p-value (sub.)	n.s.	n.s.	< 0.001
p-value (cond.)	n.s.	n.s.	n.s.
p-value (var.)	n.s.	< 0.001	< 0.009
p-value (sub*cond)	n.s.	< 0.001	n.s.
p-value (sub*var)	n.s.	n.s.	< 0.008
p-value (cond*var)	n.s.	n.s.	< 0.03

Table 4 Mean and standard deviation for the parameters modulating the speed profile for the movements of the wrist. Three-way ANOVA results for each movement type across subjects, condition and variant have been provided for each of these parameters, as well as interaction effects of the factors.

	u_x	u_y	u_z	v_x	v_y	v_z
Intrans Norm	0.36 ± 0.23	-0.36 ± 0.24	0.18 ± 0.19	0.87 ± 0.71	-0.92 ± 0.65	0.96 ± 0.47
Intrans Elb	0.68 ± 0.30	-0.54 ± 0.34	0.35 ± 0.59	0.32 ± 0.96	-0.24 ± 0.87	1.78 ± 0.60
Trans Norm	0.41 ± 0.26	-0.39 ± 0.26	0.10 ± 0.18	1.28 ± 0.56	-0.88 ± 0.72	0.77 ± 0.51
Trans Elb	0.73 ± 0.60	-0.73 ± 0.45	0.02 ± 0.50	0.71 ± 0.78	-0.14 ± 1.13	1.76 ± 1.00
p-value (sub.)	< 0.001	< 0.001	< 0.001	< 0.001	< 0.001	< 0.001
p-value (cond.)	n.s.	< 0.001	< 0.001	< 0.001	n.s.	n.s.
p-value (var.)	< 0.001	< 0.001	n.s.	< 0.001	< 0.001	< 0.001
p-value (sub*cond)	< 0.001	< 0.005	< 0.03	n.s.	n.s.	< 0.001
p-value (sub*var)	< 0.001	< 0.001	< 0.001	< 0.001	< 0.001	< 0.001
p-value (cond*var)	n.s.	< 0.02	< 0.001	n.s.	n.s.	n.s.

Table 5 Mean and standard deviation for each parameter u and v of the model describing the force at the start and end of the movements of the wrist. Three-way ANOVA results for each movement type across subjects, condition and variant have been provided for each of these parameters, as well as interaction effects of the factors.

	β	γ	δ
Intrans Norm	1.73 ± 0.66	0.011 ± 0.005	1.34 ± 0.14
Intrans Elb	1.55 ± 0.85	0.009 ± 0.004	1.22 ± 0.34
Trans Norm	1.64 ± 0.60	0.011 ± 0.005	1.33 ± 0.18
Trans Elb	1.40 ± 0.79	0.010 ± 0.005	1.27 ± 0.28
p-value (sub.)	< 0.001	< 0.001	< 0.001
p-value (cond.)	n.s.	n.s.	n.s.
p-value (var.)	< 0.03	n.s.	< 0.004
p-value (sub*cond)	n.s.	n.s.	n.s.
p-value (sub*var)	n.s.	< 0.002	< 0.001
p-value (cond*var)	n.s.	n.s.	n.s.

Table 6 Mean and standard deviation for the parameters modulating the speed profile for the movements of the elbow. Three-way ANOVA results for each movement type across subjects, condition and variant have been provided for each of these parameters, as well as interaction effects of the factors.

	u_x	u_y	u_z	v_x	v_y	v_z
Intrans Norm	0.61 ± 0.17	-0.34 ± 0.22	-0.08 ± 0.14	1.50 ± 0.033	-0.59 ± 0.39	-0.09 ± 0.58
Intrans Elb	0.85 ± 0.36	-0.63 ± 0.22	-0.06 ± 0.38	1.48 ± 0.43	-0.25 ± 0.64	1.09 ± 0.93
Trans Norm	0.50 ± 0.11	-0.38 ± 0.23	-0.03 ± 0.17	1.57 ± 0.38	-0.82 ± 0.55	-0.16 ± 0.35
Trans Elb	0.73 ± 0.36	-0.67 ± 0.27	-0.21 ± 0.29	1.44 ± 0.44	-0.56 ± 0.70	1.21 ± 1.05
p-value (sub.)	< 0.001	< 0.001	< 0.001	< 0.001	< 0.001	< 0.001
p-value (cond.)	< 0.001	n.s.	< 0.03	n.s.	< 0.001	n.s.
p-value (var.)	< 0.001	< 0.001	< 0.001	n.s.	< 0.001	< 0.001
p-value (sub*cond)	n.s.	n.s.	n.s.	n.s.	n.s.	n.s.
p-value (sub*var)	< 0.001	< 0.001	< 0.001	< 0.002	< 0.0001	< 0.001
p-value (cond*var)	n.s.	n.s.	< 0.001	n.s.	n.s.	n.s.

Table 7 Mean and standard deviation for each parameter u and v of the model describing the force at the start and end of the movements of the elbow. Three-way ANOVA results for each movement type across subjects, condition and variant have been provided for each of these parameters, as well as interaction effects of the factors.

		u_x	u_y	u_z	v_x	v_y	v_z
Sub. 1	Intrans Norm	-0.07 ± 0.20	0.03 ± 0.13	0.14 ± 0.39	-0.14 ± 0.41	0.11 ± 0.30	0.17 ± 0.45
	Trans Norm	-0.05 ± 0.15	0.02 ± 0.05	0.14 ± 0.39	-0.13 ± 0.40	0.15 ± 0.42	0.16 ± 0.44
Sub. 2	Intrans Norm	-0.05 ± 0.13	0.04 ± 0.13	0.10 ± 0.31	-0.14 ± 0.38	0.15 ± 0.43	0.17 ± 0.44
	Trans Norm	-0.06 ± 0.19	0.04 ± 0.13	0.17 ± 0.46	-0.15 ± 0.44	0.09 ± 0.25	0.16 ± 0.44
Sub. 3	Intrans Norm	-0.07 ± 0.20	0.04 ± 0.11	-0.03 ± 0.11	-0.14 ± 0.45	0.13 ± 0.35	0.17 ± 0.44
	Trans Norm	-0.09 ± 0.25	0.01 ± 0.03	0.05 ± 0.17	-0.11 ± 0.38	0.13 ± 0.35	0.21 ± 0.59
Sub. 4	Intrans Norm	-0.07 ± 0.20	0.03 ± 0.10	0.04 ± 0.17	-0.15 ± 0.42	0.20 ± 0.55	0.16 ± 0.44
	Trans Norm	-0.07 ± 0.20	0.03 ± 0.12	0.10 ± 0.30	-0.13 ± 0.40	0.18 ± 0.50	0.16 ± 0.45
Sub. 5	Intrans Norm	-0.05 ± 0.13	0.01 ± 0.04	0.21 ± 0.55	-0.13 ± 0.35	0.12 ± 0.33	0.17 ± 0.44
	Trans Norm	-0.05 ± 0.14	-0.01 ± 0.04	0.23 ± 0.61	-0.14 ± 0.39	0.04 ± 0.17	0.20 ± 0.54
Sub. 6	Intrans Norm	-0.03 ± 0.09	0.00 ± 0.01	0.14 ± 0.43	-0.12 ± 0.35	0.04 ± 0.13	0.14 ± 0.41
	Trans Norm	-0.02 ± 0.05	0.01 ± 0.03	0.23 ± 0.62	-0.10 ± 0.28	0.05 ± 0.16	0.17 ± 0.45
Sub. 7	Intrans Norm	-0.01 ± 0.08	0.01 ± 0.05	0.12 ± 0.45	0.03 ± 0.26	0.16 ± 0.44	0.20 ± 0.55
	Trans Norm	-0.03 ± 0.11	0.00 ± 0.01	0.18 ± 0.52	0.07 ± 0.21	0.06 ± 0.23	0.20 ± 0.55
Sub. 8	Intrans Norm	-0.01 ± 0.04	0.01 ± 0.04	0.15 ± 0.41	-0.15 ± 0.41	0.04 ± 0.12	0.14 ± 0.38
	Trans Norm	-0.03 ± 0.09	0.01 ± 0.03	0.18 ± 0.49	-0.18 ± 0.48	0.08 ± 0.24	0.16 ± 0.43

Table 8 Mean and standard deviation for each parameter of the model describing the force for Intrans Norm and Trans Norm movements of the wrist for each subject respectively.

		β	γ	δ
Sub. 1	Intrans Norm	0.43 ± 1.60	0.001 ± 0.002	0.04 ± 0.12
	Trans Norm	0.21 ± 0.57	0.002 ± 0.005	0.04 ± 0.12
Sub. 2	Intrans Norm	0.22 ± 0.58	0.002 ± 0.005	0.06 ± 0.17
	Trans Norm	0.36 ± 1.24	0.001 ± 0.003	0.05 ± 0.15
Sub. 3	Intrans Norm	0.25 ± 0.68	0.001 ± 0.004	0.03 ± 0.08
	Trans Norm	0.42 ± 1.46	0.002 ± 0.006	0.03 ± 0.08
Sub. 4	Intrans Norm	0.36 ± 1.57	0.001 ± 0.003	0.03 ± 0.09
	Trans Norm	0.29 ± 1.13	0.001 ± 0.004	0.02 ± 0.08
Sub. 5	Intrans Norm	0.22 ± 0.59	0.001 ± 0.003	0.04 ± 0.12
	Trans Norm	0.32 ± 1.01	0.001 ± 0.002	0.07 ± 0.20
Sub. 6	Intrans Norm	0.15 ± 0.46	0.001 ± 0.004	0.03 ± 0.08
	Trans Norm	0.23 ± 0.64	0.001 ± 0.002	0.04 ± 0.11
Sub. 7	Intrans Norm	0.24 ± 0.94	0.001 ± 0.004	0.10 ± 0.29
	Trans Norm	0.29 ± 0.94	0.002 ± 0.006	0.12 ± 0.31
Sub. 8	Intrans Norm	0.17 ± 0.45	0.001 ± 0.003	0.03 ± 0.09
	Trans Norm	0.21 ± 0.57	0.001 ± 0.003	0.04 ± 0.11

Table 9 Mean and standard deviation for each parameter of the model describing the force for Intrans Norm and Trans Norm movements of the wrist and for each subject respectively.

national Conference of the IEEE Engineering in Medicine and Biology Society

Koshland G, Galloway J, Nevoret-Bell C (2000) Control of the wrist in three-joint arm movements to multiple directions in the horizontal plane. Journal of Neurophysiology

Magescas F, Prablanc C (2006) A joint-centered model accounts for movement curvature and spatial variability. Neuroscience Letters

Meyer D, Abrams R, Kornblum S, Wright C, Smith J (1988) Optimality in human motor performance: Ideal control of rapid aimed movements. Psychological Review

Moran D, Schwartz A (1999) Motor cortical representation of speed and direction during reaching. Journal of Neurophysiology

Morasso P (1981) Spatial control of arm movements. Experimental Brain Research

Murphy J, Kwan H, MacKay W, Wong Y (1982) Precentral unit activity correlated with angular components of a compound arm movement. Brain Research

Neter J, Kutner M, Nachtsheim C, Wasserman W (1996) Applied linear statistical models. Irwin

Ogihara N, Yamazaki N (1999) Generation of human reaching movement using a recurrent neural network model. In: Proc. of the International Conference on Systems, Man, and Cybernetics

Oldfield R (1971) The assessment and analysis of handedness: the edinburgh inventory. Neuropsychologia

Plamondon R, Alimi A (1997) Speed/accuracy trade-offs in target-directed movements. Behavioral and Brain Sciences

Reina G, Moran D, Schwartz A (2001) On the relationship between joint angular velocity and motor cortical discharge during reaching. Journal of Neurophysiology

Rosenbaum D, Loukopoulos L, Meulenboek R, Vaughan J, Engelbrecht S (1995) Planning reaches by evaluating stored postures. Psychological Review

Scott S, Sergio L, Kalaskan J (1997) Reaching movements with similar hand paths but different arm orientations. ii. activity of individual cells in dorsal premotor cortex and parietal area 5. Journal of Neurophysiology

Torres E, Zipser D (2002) Reaching to grasp with a multijointed arm. i. computational model. The American Physiological Society

Uno Y, Kawato M, Suzuki R (1989) Formation and control of optimal trajectory in human multijoint arm movement. Biological Cybernetics

Wang W, Chan S, Heldman D, Moran D (2006) Motor cortical representation of position and velocity during reaching. Journal of Neurophysiology

Biol Cybern (2009) 100:331–350
DOI 10.1007/s00422-009-0300-2

ORIGINAL PAPER

Biological
Cybernetics

Movement curvature planning through force field internal models

Biljana Petreska · Aude Billard

Received: 18 September 2008 / Accepted: 16 February 2009 / Published online: 21 April 2009

Abstract Human motion studies have focused primarily on modeling straight point-to-point reaching movements. However, many goal-directed reaching movements, such as movements directed towards oneself, are not straight but rather follow highly curved trajectories. These movements are particularly interesting to study since they are essential in our everyday life, appear early in development and are routinely used to assess movement deficits following brain lesions. We argue that curved and straight-line reaching movements are generated by a unique neural controller and that the observed curvature of the movement is the result of an active control strategy that follows the geometry of one's body, for instance to avoid trajectories that would hit the body or yield postures close to the joint limits. We present a mathematical model that accounts for such an active control strategy and show that the model reproduces with high accuracy the kinematic features of human data during unconstrained reaching movements directed toward the head. The model consists of a nonlinear dynamical system with a single stable attractor at the target. Embodiment-related task constraints are expressed as a force field that acts on the dynamical system. Finally, we discuss the biological plausibility and neural correlates of the model's parameters and suggest that embodiment should be considered as a main cause for movement trajectory curvature.

Keywords Motor control · Neural control of movement · Dynamical systems · Computational model · Goal-directed reaching movements

B. Petreska · A. Billard (✉)
Learning Algorithms and Systems Laboratory,
Ecole Polytechnique Fédérale de Lausanne,
EPFL-STI-I2S-LASA, Station 9,
1015 Lausanne, Switzerland
e-mail: biljana.petreska@a3.epfl.ch

1 Introduction

The vast majority of motor control studies have focused on highly constrained reaching movements, limiting the movements to a two-dimensional plane, and in particular to the frontal plane. These constraints are meant to ensure the reproducibility and controllability of the task. They have led to the observation of so-called "quasi-straight" reaching movements with a stereotyped single-peaked, bell-shaped velocity profile (Morasso 1981; Flash and Hogan 1985). The gentle curvature responsible for the term "quasi" has proved hard to explain. Some have suggested that it is due to distortions in the visual perception of the target (Wolpert et al. 1994, 1995), which could however not explain the fact that these are also observed in congenitally blind subjects (de Graaf et al. 1994). Others have attributed the curvature of the movement to the dynamics of the arm's biomechanics, i.e., inertial and viscoelastic resistive forces (Flash 1987; Bullock and Grossberg 1988). This again could not explain the fact that the curvature persists in isometric tasks, which indicates rather that the curvature is encoded directly in the activation patterns of the muscles (Pellegrini and Flanders 1996). Another possible explanation for the curvature of arm movements is Listing's law, as the arm rotation movements were shown to roughly lie in a 2D curved surface (Liebermann et al. 2006). Importantly, when participants are instructed to generate straight paths, they produce movements much straighter than those generated spontaneously (de Graaf et al. 1994; Desmurget et al. 1997; Osu et al. 1997), which argues against the hypothesis of imperfect control (Flash and Hogan 1985). In addition, the curvature depends on the location of the target (Soechting and Lacquaniti 1981) and is systematic within trials and across subjects (Soechting and Lacquaniti 1981; Pellegrini and Flanders 1996). Curved trajectories are also more frequently observed during unconstrained movements

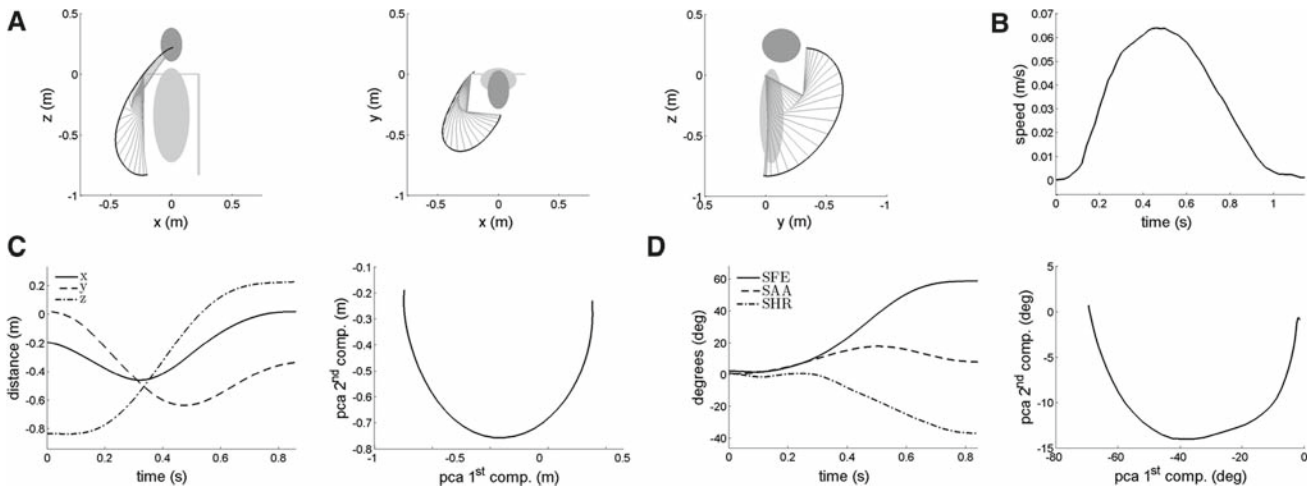

Fig. 1 An example of the curvature of an unconstrained self-oriented movement (the subject was asked to touch his nose). **a** Projections of the movement in the *xy*-, *xz*-, and *yz*-planes. **b** The velocity profile is bell-shaped and single-peaked, similarly to the velocity profiles of straight point-to-point movements. **c** The movement is curved in the extrinsic hand Cartesian space (*left*), which is best visible when projected on the first two principal components following a principal component analysis (PCA) (*right*). **d** The movement is curved also in the intrinsic joint angles space (*left*) and its two principal components (*right*). The joint angles represented here correspond to the three degrees of freedom of the shoulder: shoulder flexion–extension (*SFE*), shoulder abduction–adduction (*SAA*) and shoulder humeral rotation (*SHR*)

(Soechting and Lacquaniti 1981; Lacquaniti et al. 1986; Miall and Haggard 1995; Desmurget et al. 1997; Osu et al. 1997). Overall, the above evidence indicates that the curvature underlying human motion might be a "natural" feature of the movement, and the observed straightness an artifact of the restricted workspace.

We show in this paper that these non-linearities are particularly important when considering reaching movements directed to ourselves (see Fig. 1). Self-oriented movements are part of our daily repertoire (e.g., to eat). They are among the first to emerge in life and are likely the result of evolutionary old neural structures. Their study may thus reveal basic neural processes of motor control. For instance, electrical stimulation of the precentral and motor cortices evoked natural multijoint movements that reached to different points in space, such as for example characteristic hand-to-mouth movements (Graziano et al. 2002, 2005). These movements are also routinely used in neurological examinations to test and diagnose various movement deficits following brain lesion (De Renzi and Lucchelli 1988; Goldenberg and Hagmann 1997; Petreska et al. 2007), which directly inspired the stimuli used in our study. All in all, the study of reaching movements toward oneself is particularly interesting from both a behavioral and a neurological perspective.

We will argue that movement curvature is planned by the central nervous system (CNS) and takes into account the geometry of the body. The idea that embodiment can be encapsulated in the control system itself is in line with our earlier observation that differences in the kinematic features of reaching movements in macaques and humans could be related to the biomechanical properties of the macaques' and humans' shoulder joints (Christel and Billard 2002). Importantly, the model proposed here is not limited to self-oriented movements and can be applied to any point-to-point reaching movement such as for example reaching to targets in the extrapersonal frontal workspace.

2 Computational approach

Modeling studies are particularly useful for distinguishing among all of the plausible mechanisms to encode movements, as long as their predictions are tested and validated against empirical behavioral or neurophysiological data.

However, existing models are unsuccessful at reproducing the curvature of natural human movements (Admiraal et al. 2004), up to several exceptions (Torres and Zipser 2002; Biess et al. 2007; Guigon et al. 2007). For instance, while the so-called 2/3 power law (Lacquaniti et al. 1983) could account well for the curvature observed during handwriting and drawing motions, it was unsuccessful at explaining the curvature of reaching movements in the 3-dimensional space (Schaal and Sternad 2001), including the movements we consider in this paper as shown on Fig. 2. Furthermore, the minimum work model (Soechting et al. 1995) successfully reproduces the final joint postures of pointing movements starting from different initial joint postures, but does not explain the time dependency across joint trajectories. A kinematic model that intrinsically constrains the arm joints according to Listing's law (i.e., such that the arm rotation vectors lie in a 2-dimensional surface) was partially successful at describing the experimental data (Liebermann et al.

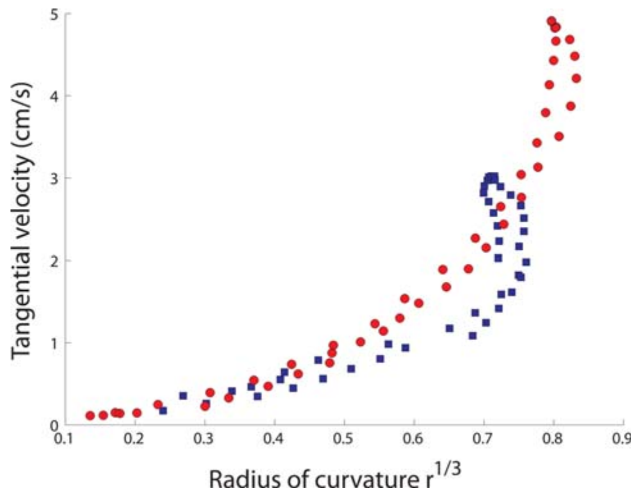

Fig. 2 Two examples of unconstrained self-oriented movements where the 2/3 power law was degraded. The tangential velocity versus radius of curvature to the power 1/3 is shown. The subject was asked to touch his nose (*circles*) or to touch his left ear (*squares*)

2006). The minimum hand jerk[1] model (Flash and Hogan 1985) maximizes the smoothness of the hand trajectory in the extrinsic space. The result is a straight-line trajectory, whereas curved trajectories are obtained by specifying via-points (e.g. for avoiding obstacles). However, it predicts a bimodal velocity profile which is at odds with the experimental data (Atkeson and Hollerbach 1985). Later it was suggested that the hand trajectory is the result of a compromise between planning a straight line in the task space and planning a straight line in the joint space (Cruse and Brüwer 1987; Okadome and Honda 1999; Hersch and Billard 2007). Such hybrid computations offer numerous advantages for controlling 3-dimensional reaching movements, such as avoiding singularities and avoiding hitting the joint limits (Hersch and Billard 2007). Unfortunately there is currently no direct neurophysiological evidence in support of such a control strategy. It has also been proposed that arm movements are controlled by minimizing the derivative of joint torques (Uno et al. 1989; Nakano et al. 1999; Wada et al. 2001). However, this model overestimates the magnitude of curvature of pointing movements (Biess et al. 2007). In Torres and Zipser (2002), the hand path is computed in the intrinsic joint angles space by minimizing an energy-like quantity, giving realistic predictions for curved paths. However, this model assumes a separate processing for the spatial and temporal dynamics of motion and displays some imprecisions for movements similar to those addressed here. The model by Biess et al. (2007) computes a geometrical joint angles geodesic path with respect to a kinetic energy metric in the Riemannian configuration space and subsequently minimizes the squared jerk along this path. This model also treats the spatial and temporal dimensions separately and predicts identical path trajectories for different speeds. We find it difficult to evaluate how well this model would predict highly curved reaching movements as the pointing movements addressed in the study were quasi-straight, but we could observe that the model has difficulties with reproducing mixed curvatures (i.e., movements that deviate first to one side and then to the other side of the idealized straight trajectory). Another class of reaching models are stochastic models that take into account the noise inherent to the motor system. It has been consistently observed that the standard deviation of neuromotor commands increases with its mean (Sutton and Sykes 1967; Schmidt et al. 1979; Clamman 1969; Matthews 1996; St-Amant et al. 1998; Clancy and Hogan 1999; Osu et al. 2004). In line with this evidence, it was suggested that the brain minimizes the variance of the final arm position in the presence of such signal-dependent motor noise (Harris and Wolpert 1998; Hamilton et al. 2002). Even though this model succeeds at reproducing the curvature of 2-dimensional reaching movements, it does not specify which control laws generate these movements. In Todorov and Jordan (2002), an optimal feedback theory of motor control is proposed, in which the variability of the movement is distributed optimally among different degrees of freedom that do not interfere with the task goal. This qualitative model is appropriate for explaining the variability observed in reaching movements, it is however imprecise in its prediction of the curvature of movements. This is partly due to the determination of the appropriate cost function to optimize. This performance criterion is chosen arbitrarily and varies with the task. Another model based on the optimal feedback control theory was successful at reproducing the joint and hand trajectories of 3-dimensional movements (Guigon et al. 2007), but the authors admit that the movements reproduced are rather stereotyped. For example the model does not account for nonsymmetric velocity profiles or avoidance of extreme joint limits.

While it has been suggested that two different control strategies underlie straight and curved reaching movements (Desmurget et al. 1997; Moran and Schwartz 1999), we argue that these two types of movements are generated by a unique adaptive control mechanism. While none of the existing models offers a satisfactory solution for modeling the highly variable curvature of human movements, here we propose a dynamical model that accounts for both gently and highly curved hand trajectories, consistent with recent neurophysiological findings. First, unlike many of the models above, our model is closed-loop. Closed-loop control takes into account the uncertainty of the "real-world" and allows intelligent online corrections as well as robust responses to perturbations, rather than "playing a prerecorded tape" (Todorov 2004). Such an approach is in agreement with the observa-

[1] The jerk corresponds to the derivative of the acceleration and is a measure of the smoothness of the trajectory.

tion that the CNS is able to estimate and anticipate the state of the limb. This is achieved by integrating delayed sensory input and motor output through afferent and efferent internal feedback loops (Desmurget and Grafton 2000). The state information is used to continuously update the motor commands, which is likely to occur in the posterior parietal cortex and cerebellum.

Our model also takes advantage of the signal-dependant neuromotor noise mentioned earlier, which may be responsible for the speed-accuracy trade-off known as Fitts' law (Fitts 1954) and trail-to-trial variability (Todorov 2004). Finally, our model hypothesizes that the curvature of the hand trajectories is not an undesirable noise on otherwise perfect straight-line reaching movements. Rather, it is necessary and planned as such by the CNS in order to, for example, avoid impossible trajectories that go through the body and uncomfortable joint limit postures.

3 Model description

Our work was driven by the assumption that (a) a unique controller underlies both straight and curved reaching movements, and (b) that this controller is such that all the variables can be accounted for by known neurophysiological processes. Thus, to start with, we considered the vector integration to endpoint (VITE) model for point-to-point reaching (Bullock and Grossberg 1988) that accounts for typical kinematic features of human reaching movements such as bell-shaped velocity profiles and speed-accuracy trade-off. The model has been used to explain control in both hand extrinsic and joints intrinsic spaces (Ajemian et al. 2001; Hersch and Billard 2007). Most importantly, the dynamics of the VITE model's response displays a profile of activity similar to that of populations of neurons in the primate's brain. In particular, the model could account for these neurons' sensitivity to change in the velocity of the movement and for the latency of activity at the movement onset (Bullock et al. 1998). The VITE model, however, suffers from a major restriction: it can generate only *straight* movements.[2] Next, we describe the VITE model and give a formal definition of our extension that accounts for curved reaching movements.

[2] An extension of the VITE model has been proposed to account for highly curved handwriting movements (Bullock et al. 1993; Paine et al. 2004), where three coupled VITE models control the displacement of the hand in a 2-dimensional plane and the rotation of the wrist. The curvature results from the coupling between the three models and the fact that each model is initiated with a slight delay at onset. This approach is not optimal for modeling simple point-to-point reaching movements as it necessitates the characterization of a sequence of multiple arbitrary targets, one for each change in the curvature.

3.1 The original VITE model

The original VITE model is a dynamic controller that at each point in time reduces the distance between the estimated and desired states of the controlled variable. First, it computes the desired movement acceleration based on the difference between the present and endpoint vectors. Second, this acceleration is integrated and primed with a faster-than-linear time-dependent "go signal" to specify the desired speed, which is the control signal sent to the muscle motoneurons. This priming signal is essential for the obtention of a bell-shaped velocity profile.

In its complete form the VITE model succeeds for example at: maintaining accurate proprioception while controlling voluntary reaches to spatial targets, maintaining postures despite perturbations, complying with an imposed movement, exerting force against obstacles, compensating for static and inertial loads and reproducing muscle vibration effects (Cisek et al. 1998). For simplicity, we only use the concise form of the model presented in Bullock and Grossberg (1988). For a description of the original VITE model please see the Appendix.

3.2 Modification of the original VITE model

Our modified VITE system is governed by a non-linear and noisy spring-damper system given by:

$$\ddot{\mathbf{x}}(t) = \overbrace{-\alpha\dot{\mathbf{x}}(t)}^{\text{damping factor}} + \overbrace{\beta\mathbf{g}(\mathbf{x}^*(t) - \mathbf{x}(t) + \boldsymbol{\eta})}^{\text{noisy endpoint attractor}}\,\mathbf{u}(t) \tag{1}$$

The first term is a damping factor proportional to the speed $\dot{\mathbf{x}}(t)$ of the end-effector that prevents the system from oscillating too importantly. The second term corresponds to an elastic force that drives the end-effector from its actual position $\mathbf{x}(t)$ toward the desired target position $\mathbf{x}^*(t)$. Note that the desired position is written as a function of time in order to emphasize the ability of the system to track the target in real time without any additional computation (as a result the system is robust to perturbations of the target position). α is a time constant set to 50. $\beta \in \mathbb{R}^+$ determines the amplitude of the speed at which the system moves globally (increasing β would result in a higher velocity peak and shorter movement duration, see Fig. 3a). $\mathbf{g}$ is a nonlinear function that modulates the dynamics of the system so that it presents a typical bell-shaped velocity profile (refer to the Appendix for the exact form of $\mathbf{g}$). Finally $\boldsymbol{\eta}$ is a multiplicative gaussian noise with zero mean and standard deviation proportional (by a factor of 0.005) to the distance between the actual and desired end-effector positions, namely $|\mathbf{x}^*(t) - \mathbf{x}(t)|$. This noise factor is necessary to initiate the movement and to account for the trial-to-trial variability at the onset of movement (see the Appendix).

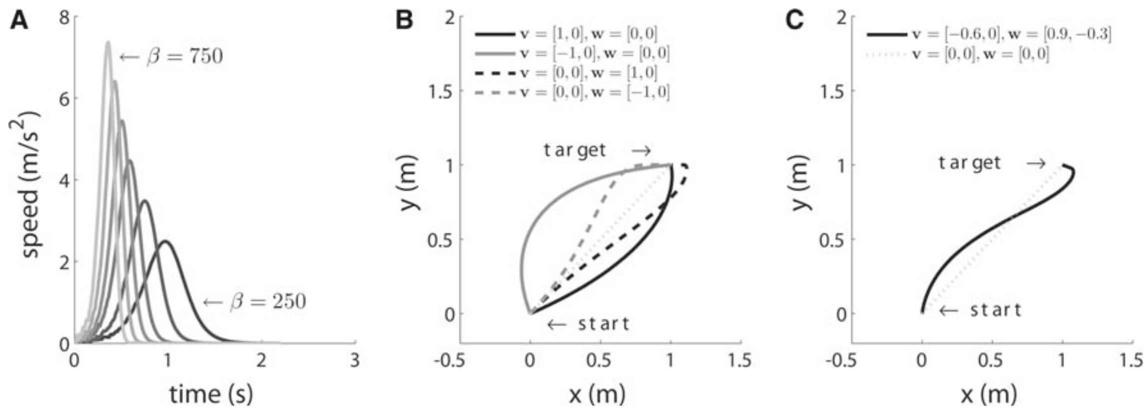

Fig. 3 **a** Effect of gradually increasing the parameter β of the modified VITE model (see Eq. 1) on the velocity profile of the movement. Higher β values increase the velocity peak and shorten the movement duration. **b** Behavior of the extended F2REACH model (see Eq. 2) under different repulsive forces **v** and **w**, for illustrative purpose the forces shown are applied only on the horizontal dimension. The forces are modulated such that **v** affects mostly the beginning of the movement and **w** mostly the end of the movement. Note that the direction of the deviation from the straight trajectory is determined by the sign of the force. **c** By combining two forces **v** and **w** of different signs one can obtain very interesting deviations that change their direction during the execution of the movement. Reference values: $\alpha = 50$, $\beta = 500$, noise was set to 0.005. Only the speed parameter β is varied throughout the simulations, the two other parameters (time constant and noise) are fixed to the given values

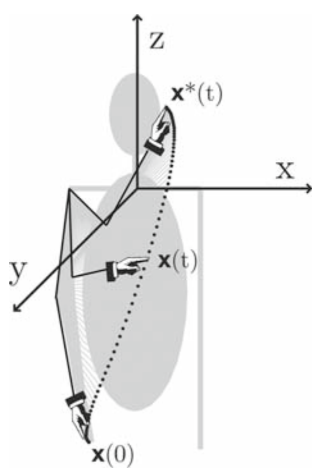

Fig. 4 Description of the task space. The hand position $\mathbf{x}(t)$ is represented in a 3-dimensional space centered on the chest, at the level of the shoulders. The input to the model consists of the initial hand position $\mathbf{x}(0)$ and final target position $\mathbf{x}^*(t)$

The above formulation makes two strong assumptions from a motor control point of view: (a) it takes as control signal the acceleration of the end-effector $\ddot{\mathbf{x}}$, expressed in an extrinsic 3-dimensional Euclidean space centered on the chest (see Fig. 4), and (b) it accounts only for a "high-level" control mechanism, in that it generates the desired end-effector kinematics, and does not account for the subsequent transformation required to control muscle activations.

Expressing the system in terms of desired acceleration is not constraining, since it is conceivable to assume that a neural population coding for the acceleration can be neurally integrated out to obtain a velocity control signal, which can in turn be integrated out to have a position control signal, see Sauser and Billard (2006). Moreover, evidence that muscle activity may be governed by a kinematic signal, such as the acceleration, velocity or position, or any combination of these, has been found in the motor cortex (Wang et al. 2007). Note that we do not address the problem of redundancy mapping between desired hand kinematics and actual muscle activations in this paper. These assumptions will be further developed in Sect. 6.1.

The above system differs from the original VITE model in two ways (see the Appendix for the original VITE formulation). First, the dynamics of the system is now governed by a *single* second order differential equation and is thus expressed in terms of the end-effector acceleration.[3] Second, we replaced the explicit time dependency of the original VITE system by introducing a bounded nonlinearity in the function **g**. In the original VITE system, this explicit dependency in time through the priming signal let the velocity of the system grow exponentially in time, which created instabilities in the case of a long lasting perturbation, and was thus biologically implausible (your arm does not start accelerating if someone holds it).

3.3 Extension of the original model: F2REACH model

To account for the movement curvature, we next introduce a new functional $\mathbf{F}(\mathbf{x}(t))$ that corresponds to a *virtual force*

[3] The original VITE system was driven by two coupled first-order differential equations. We reformulated this by writing the whole system as a second order differential equation. This allows us to relate explicitly the acceleration of the system to the force-field which we introduce in the following section.

field, which encapsulates a geometrical representation of the task constraints. This force field is modulated by the dynamics of the control signal in order to preserve the bell-shaped velocity profile:

$$\ddot{\mathbf{x}}(t) = -\alpha\dot{\mathbf{x}}(t) + \beta\mathbf{g}(\mathbf{x}^*(t) - \mathbf{x}(t) + \boldsymbol{\eta}) + \overbrace{\beta|\mathbf{g}(\mathbf{x}^*(t) - \mathbf{x}(t) + \boldsymbol{\eta})|}^{\text{modulation factor}}\,\overbrace{\mathbf{F}(\mathbf{x}(t))}^{\text{force field}} \quad (2)$$

The force field $\mathbf{F}(\mathbf{x}(t))$ assigns a vector gradient to each position in space that expresses constraints related to: (a) objects in the environment that one needs to avoid (including the subject's body), (b) dynamic properties of the human body such as inertial properties of the limb, (c) extreme joint angles limits. The contribution of each of these constraints is simply summed to result in the virtual force field. The gradient of the force field at each point in space pushes the hand away from the undesired locations.

This force field framework reconciles the dynamic and kinematic aspects as well as intrinsic and extrinsic approaches to motor planning in a very convenient way. Instead of finding a compromise across systems that would operate simultaneously in conflicting coordinates (e.g., hand position and joint angles, see Sect. 2), our system provides both dynamic (acceleration) and kinematic (speed or position) control signals, taking into account (a) a target for the motion expressed in extrinsic kinematic coordinates and (b) intrinsic dynamic motion constraints. This reconciles the observation that objects in the environment such as a table may influence the kinematic planning of the movement[4] (Brenner and Smeets 1995) and that knowledge of the arm dynamics is necessary for the kinematic planning of complex movements (Uno et al. 1989; Nakano et al. 1999; Sabes and Jordan 1997).

As the particular form taken by the force field is task and context dependent, we chose a very generic expression given by:

$$\mathbf{F}(\mathbf{x}(t)) = h(\mathbf{x}(t))\mathbf{v} + (1 - h(\mathbf{x}(t)))\mathbf{w} \quad (3)$$

where $\mathbf{v}$ and $\mathbf{w}$ are constant force vectors that push the trajectory away from the straight-line generated by the rest of the system. $\mathbf{v}$ affects primarily the *beginning* of the movement, whereas $\mathbf{w}$ affects the *end of movement* (as illustrated in Fig. 3b, c). The modulation function h that associates these two forces to different parts of the movement is given in the Appendix.

In our framework, a 3-dimensional reaching movement needs the specification of seven parameters in total: β that controls the amplitude of the velocity's peak and two 3-dimensional repulsive forces $\mathbf{v}$ and $\mathbf{w}$, where the time constant α and noise can be fixed to 50 and 0.005 respectively. We will show next that the latter two forces give a crude representation of the volume and geometry of the body around which the hand must navigate.

To conclude the description, control policies of the form of autonomous differential equations such as the one proposed here are particularly interesting, as they allow online modifications of the input variables. Thus a very nice property of our model is its robustness to external perturbations, where the model shows smooth adaptation to changes such as blocking or displacing the arm and displacing the target (simulation results not shown here).

4 Experiments

4.1 Subjects

Ten healthy subjects, five female and five male of mean age 33 ± 11 years volunteered for the study. All the participants except for two were right-handed according to the Edinburgh handedness test (Oldfield 1971). All the subjects were naive as to the purpose of the study and had no history of neurological or musculoskeletal deficits.

4.2 Procedure

The subjects were asked to perform natural reaching movements toward targets situated on their head. In order to obtain entirely natural and fully unconstrained movements, the target positions were specified verbally (for example we gave instructions such as "on the go signal touch your nose"). The subjects were left free to determine the location of the reaching target (e.g., at the tip of the nose or just above it), but they were instructed to reach to exactly the same location across one block of repetitions of the same movement. There were six target positions, shown in Fig. 5a, indexed as follows: (1) nose, (2) right ear, (3) left ear, (4) top of the head, (5) under the chin and (6) back of the head. Given that the subjects had different arm lengths and given that the targets were defined with respect to the subject's head, the length of the hand path varied importantly across subjects and movements. This was done on purpose to test the ability of the model to reproduce the generic characteristics of the movements and to account for such body variabilities, which we consider task-independent. The subjects were standing in order to limit undesirable movements of the upper body. There were no external constraints that would confine the movement range. The movements were performed with the right hand independently of the handedness of the subject, since handedness was shown not to affect spontaneous self-oriented movements (Dalby et al. 1980; Lavergne and Kimura 1987). In order to verify the gener-

[4] This type of computation is natural (and especially useful) if the movement is considered in a constantly varying environment full of external objects, instead of isolated in an artificial experimental setup.

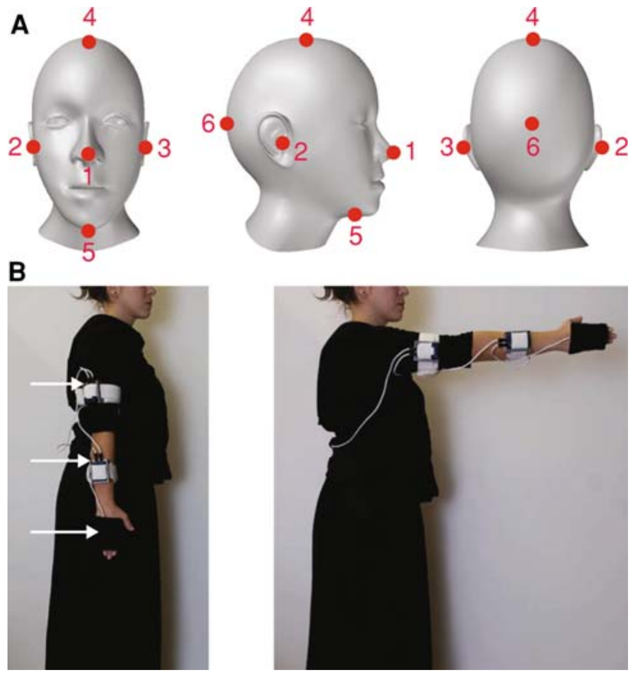

Fig. 5 **a** Target positions on the head used in our experiment. **b** Two initial conditions that yield both highly and gently curved movements. The three motion sensors are indicated with *arrows*

alization of our model over movements with different curvature levels, movements were initiated from two different locations, shown in Fig. 5b: (1) upright position with the arm extended along the body that yielded highly curved movements and (2) upright position with the arm extended in front of the body that yielded gently curved movements. Prior to each experiment, the subjects were asked to assume the same starting position, which was verified by the experimenter. The subjects had at least one trial of practice per movement to ensure that they had understood the instructions. Each movement was repeated five times in order to have a measure of its inherent variability and consequently a measure of the precision of the model's reproduction.

4.3 Data acquisition

Data was recorded using 3D inertial measurement units/ motion sensors (Xsens Technologies B.V., The Netherlands). The sensors were attached on three arm segments (the upper arm, the forearm and the hand) and were calibrated in the upright position with the arm vertical (see Fig. 5b, left). The orientation of the three arm segments during the execution of the movements was recorded at a frequency of 50 Hz.

4.4 Data analysis

All analyses were performed with custom software written in Matlab (Mathworks, Natick, MA, USA). The trajectories of each arm segment were reconstructed using the orientation matrices recorded by the inertial measurement motion sensors. We used only *unfiltered raw* values. The movements of interest were extracted using criteria such as percentage of velocity change. The samples were aligned in time so that the inter-trial Euclidean distance per movement and subject (five samples) is minimal. The movement mean and standard deviation (SD) of each trajectory for each movement type and for each subject was computed with respect to the aligned signals. We then solved numerically the original VITE and extended F2REACH models for each of the mean movements, with a time step of 20 ms. The models' parameters were fixed using 3^3 and 3^7 factorial experimental designs respectively, coupled with a local search procedure (Neter et al. 1996; Hoos and Stützle 2004).

To evaluate the predictions of the two models we measured the following Euclidean distances and deviation indices:[5] (1) *mean deviation* (MD) of the predicted hand trajectory compared to the measured hand trajectory at each point in time, (2) *mean squared error* (MSE), (3) *hand trajectory deviation index* (HTDI) defined as the ratio between the maximal distance across the modeled $\mathbf{x}^m(t)$ and real $\mathbf{x}^r(t)$ mean trajectories over the total length of the real path,

$$\mathrm{HTDI} = \frac{\max_{i=1,\ldots,N} |\mathbf{x}^m(i) - \mathbf{x}^r(i)|}{\sum_{i=1}^{N-1} |\mathbf{x}^r(i+1) - \mathbf{x}^r(i)|}$$

where N is the number of points sampled (see Fig. 6a), (4) *speed deviation index* (SDI) and finally (5) *total acceleration deviation index* (ADI), both defined in Fig. 6b. We also considered the standard deviation trajectory (SD) as a possible limit prediction (see Fig. 6c for a definition). We further assessed the *curvature index* of recorded and modeled movements, defined as the ratio between the total arc length of the hand path and the Euclidean distance that separates the initial and final positions. For example a curvature index of 1 indicates a perfectly straight path and a curvature of $\pi/2$ corresponds to a semicircular path. Finally, the *speed asymmetry index* was defined as the ratio $(S^a - S^d)/(S^a + S^d)$ where S^a is the distance traveled up to the time at which the velocity is maximal (referred to as the acceleration phase) and S^d the distance traveled from the time at which the velocity is maximal until the end of the movement (deceleration phase). An additional measure of the precision of the original VITE and extended F2REACH models is the percentage of trajectory points predicted by the models that are comprised within the volumes defined by 1 and 2 SD away from the recorded mean trajectory (per subject and movement type, established over five repetitions of the movement, see Fig. 6c). This measure accounts for the variability inherent to goal-directed reaching

[5] The deviation indices are adapted from Nakano et al. (1999) and Biess et al. (2007).

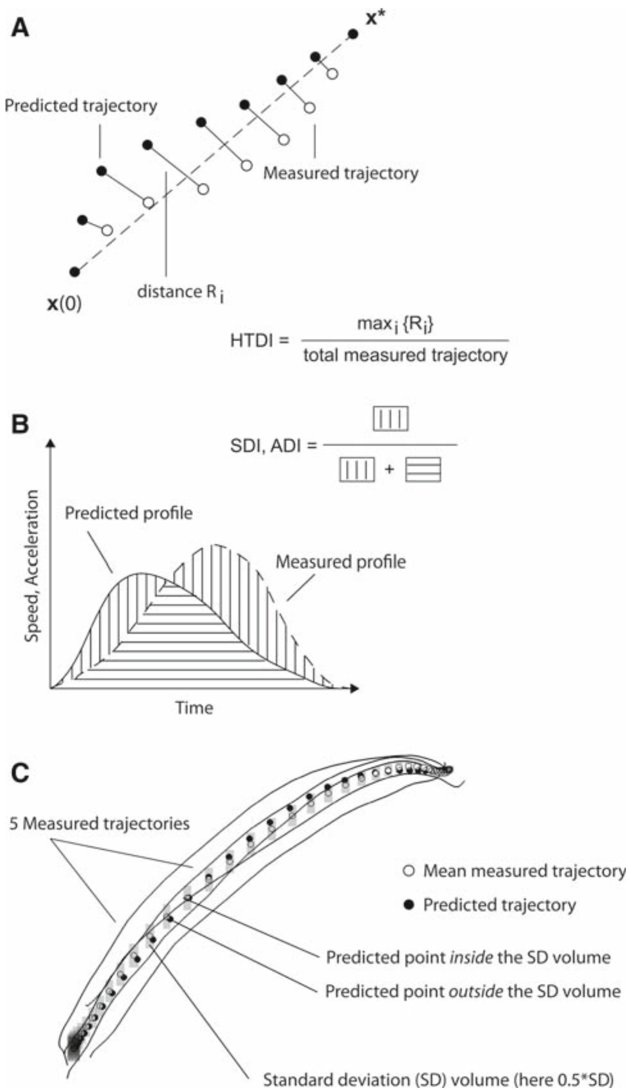

Fig. 6 Definitions of error measures. **a** The hand trajectory deviation index (*HDTI*) of measured and predicted hand trajectories is the ratio of the maximum distance, $R = \max_{i=1,\ldots,N} R_i$, between the two trajectories matched in time over the total length of the measured path. **b** The speed deviation index (*SDI*) and total acceleration deviation index (*ADI*) are defined as the ratio of the noncommon area enclosed by the measured and predicted speed/acceleration profiles and the total enclosed area. **c** Standard deviation volumes (SD), comprised within a multiple of the standard deviation distance (computed from the mean trajectory of five movement trials per subject and movement type) at every point of the movement trajectory. A SD trajectory would follow the corresponding corners of these volumes. We consider that a point was well predicted if it is contained inside the SD volume of its measured counterpart, thus enforcing a higher precision at points with very low variability

movements (Harris and Wolpert 1998; Todorov and Jordan 2002) and penalizes imprecision in parts where the variance of the movement is minimal. For example, the subjects were more consistent in the vicinity of the initial and target positions.

5 Results

In this section we report on a systematic assessment of how well the original VITE and our extended F2REACH models account for the kinematics of the recorded human movements. We also discuss the biological plausibility of our model's parameters. Finally, we conduct a stability analysis of the F2REACH model and define conditions under which the target is a stable attractor of the model and therefore guaranteed to be reached.

5.1 Observed data statistics

We first assessed the general characteristics of the recorded movements (summarized in Table 1). The movements addressed had large spatial extent (mean path length of 1.23 m) with significantly longer path lengths in the first experimental condition (see Fig. 5b) when compared to the second experimental condition (mean path lengths of 1.7 and 0.95 m respectively). Movements in the first condition lasted longer with mean durations of 1.3 and 1 s, respectively. Most importantly, the movements in the first condition were significantly more curved with a mean curvature index of 1.59 compared to 1.21 in the second condition. In addition, the curvature indices of the recorded movements were distributed homogeneously between quasi-straight (<1.1) and highly curved (>2).

We expected to see substantial trial-to-trial fluctuations due to noise of the motor system (Todorov and Jordan 2002), which motivated us to model the mean trajectory of the movement rather than the separate trials. We believe that the mean movement captures the intrinsic nature of the movement, which is task-relevant and free of noise. An example of the inherent variability across trials per subject and movement type is shown in Fig. 7a. Figure 7b shows that the inter-subject variability (attributed to the difference in embodiment of the subjects) is much more important.

5.2 Comparison between the observed and modeled data

Here we assess how well the original and extended models reproduce the human data. The mean movement trajectories were simulated with both the original VITE and our extended F2REACH models. Typical examples of measured and predicted hand path trajectories are given in Fig. 8. The first row in each example shows the five hand trajectories of the movement projected in the *xy*-, *xz*- and *yz*-planes relative to a schematized humanoid. The second row shows the projections of the mean recorded trajectory and generated model trajectories. The subject's trials are represented with light grey lines and show the inherent variability of the movement. The third row shows the *x*-, *y*- and *z*-components of the hand trajectories with respect to time in order to show the quality of the model predictions at the temporal level.

Table 1 Path length, duration and curvature index of the movements in the two experimental conditions (see Fig. 5b)

	Condition 1	Condition 2	2 Conditions	10 Subjects
			P value	*P* value
Path length (m)	1.70 ± 0.32	0.95 ± 0.18	<0.001	NS
Duration (s)	1.28 ± 0.26	0.97 ± 0.20	<0.001	<0.001
Curvature index	1.59 ± 0.22	1.21 ± 0.11	<0.001	NS

We also give one-way ANOVA results for the initial condition and subject effects on these variables. The movements in condition 1 were significantly longer in time and space and significantly more curved when compared to the movements in condition 2. The recorded movements differed significantly across subjects only in their duration

Table 2 Mean deviation (MD), mean squared error (MSE) and mean deviation indices (see Fig. 6) for the trajectory (HTDI), speed (SDI) and acceleration (ADI) (± standard deviation) of the hand as predicted by the extended F2REACH and original VITE models

	F2REACH model	SD	VITE model
MD (mm)	18.85 ± 8.10	35.67 ± 11.63	132 ± 71
MSE (cm^2)	5.62 ± 5.34	15.93 ± 10.61	431 ± 413
HTDI	0.031 ± 0.010	0.04 ± 0.02	0.25 ± 0.06
SDI	0.11 ± 0.03	0.50 ± 0.11	0.29 ± 0.12
ADI	0.38 ± 0.07	0.60 ± 0.08	0.51 ± 0.13

We also consider the trajectory comprised within one standard deviation (SD) from the mean trajectory (per subject and movement type, computed as described in Fig. 6c) as an indication for the limit prediction that would be acceptable for a model. This SD trajectory represents the inherent variability of the movement. One-way ANOVAs performed on the error measures of the extended F2REACH model show that the effect of the subject performing the movement was not significant and that the movements in the second initial condition, i.e., movements with lower curvature, tended to be slightly better predicted (MD and MSE only)

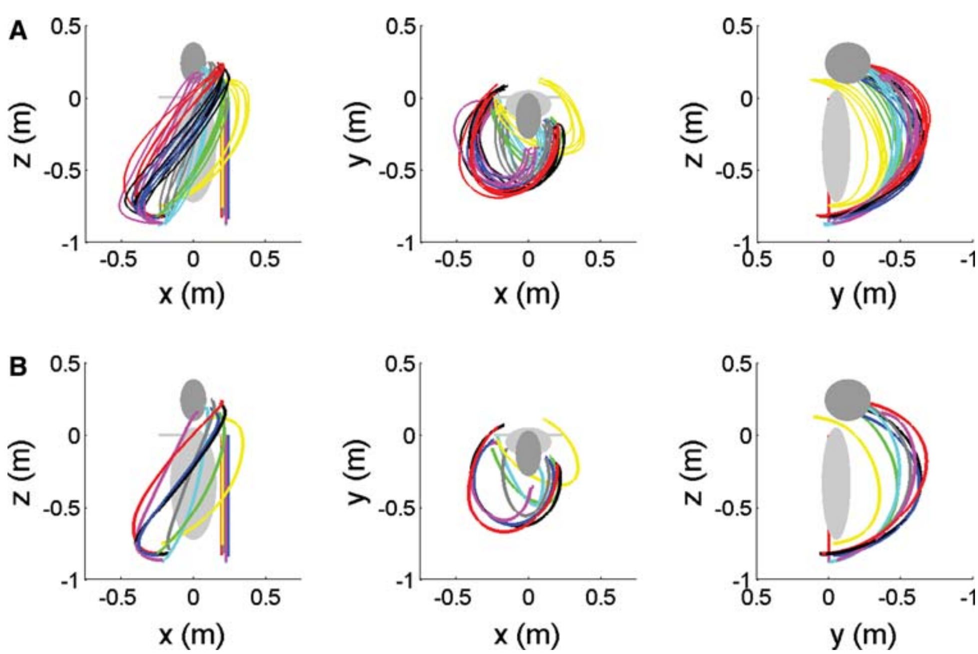

Fig. 7 Trajectories of the hand for ten subjects performing five repetitions of the same movement, reaching to the left ear (movement 3) with the right arm in condition 1 (see Fig. 5). The hand trajectories are shown relative to a schematized humanoid and the color refers to the same subject. **a** All the movement trajectories are shown in order to emphasize the movement's inherent variability. Note that this intra-subject variability is lower than the inter-subject variability, i.e., the hand trajectories of one subject are consistent when compared to those of the other subjects. **b** Only the mean movements are shown. The inter-subject variability can be partially attributed to differences in the subjects' arm lengths and shoulder positions (see *color-coded arms*)

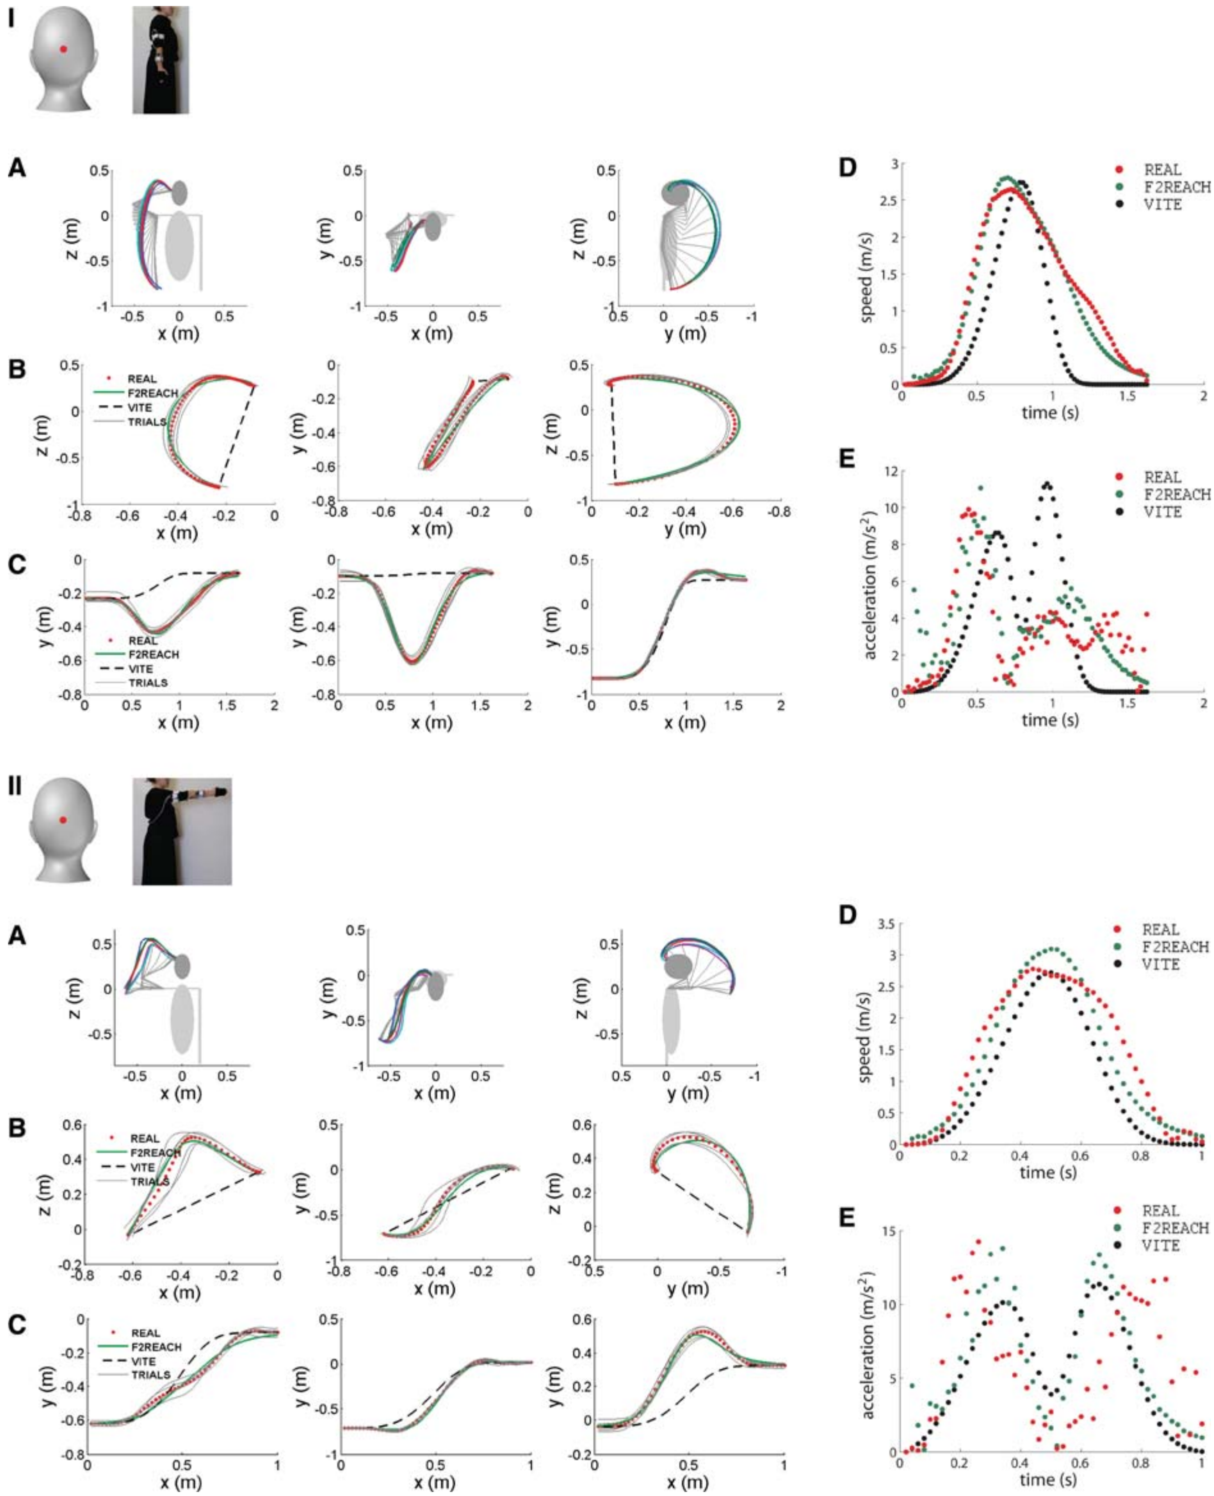

Fig. 8 Two examples of typical movements. The recorded human data is shown with points that respect the sampling rate, the original VITE model is shown with a dashed line and our extended F2REACH model with a plain line. *I*, The subject reaches for the back of the head (movement 6) with as initial condition the right arm extended along the body (condition 1). *II* The subject reaches for the back of the head (movement 6) with as initial condition the right arm extended in front of the body (condition 2). **a** The five recorded hand trajectories of the movement projected in the *xy*-, *xz*- and *yz*-planes and shown relative to a schematized humanoid. **b** The measured and predicted mean movement trajectories projected in the *xy*-, *xz*- and *yz*-planes. The light grey trajectories are the five trials and reflect the intra-subject variability per movement type. **c** The *x*-, *y*- and *z*-components of the measured and predicted mean movement trajectories shown with respect to time. **d** The measured and predicted speed profiles of the movement. **e** The measured and predicted total acceleration profiles of the movement

Table 3 Measured (M) velocity peak amplitude and peak time, asymmetry and curvature indices ($\pm$ standard deviation) compared to those predicted by the extended (F2REACH) and original (VITE) models

	Measured (M)	F2REACH model	M versus F2REACH P value	VITE model	M versus VITE P value
Velocity peak amplitude (m/s)	2.26 ± 0.68	2.37 ± 0.69	<0.05	8.22 ± 59.31	NS
Velocity peak time (s)	0.50 ± 0.14	0.53 ± 0.13	NS	0.59 ± 0.19	<0.001
Asymmetry index	-0.08 ± 0.15	-0.10 ± 0.11	NS	0.03 ± 0.06	<0.001
Curvature index	1.40 ± 0.26	1.36 ± 0.23	NS	1.03 ± 0.22	<0.001

There were no significant differences between the measured and extended model variables (with the exception of a small difference in the velocity peak amplitude), whereas significant differences were found between the measured and original model for three of the four variables addressed (all except for the velocity peak amplitude)

Finally, on the right we show the measured and predicted speed and acceleration profiles. One can see that, unlike the original VITE model, the F2REACH model is generally in very good agreement with the human data.

We systematically evaluated the predictions of the original VITE and extended F2REACH with several Euclidian distances and deviation indices defined in Sect. 4.4. The results are summarized in Table 2 and show that our model is highly precise at reproducing the kinematics of the recorded movements. The deviation indices are much smaller, generally on a different order of magnitude than those from the SD trajectory and always smaller than the original VITE model. The mean deviation was less than 2 cm for movements of average path length superior to 1 m.

We performed one-way ANOVAs for the extended model using, as dependent data, the different error measures defined in the preceding paragraph. The results show that, regardless of the error measures used, we did not find an effect of the subject executing the movements ($P > 0.05$, with the exception of two subjects for the HTDI and ADI deviation indices). This indicates that our model performed equally well across the ten subjects. A significant effect ($P < 0.001$) was observed for the two experimental conditions (see Fig. 5) for the mean deviation (MD), mean square error (MSE) and speed deviation index (SDI) suggesting that the model is better at predicting low rather than high curvatures. This result is not very surprising since the force field in our model is parameterized with two constant forces, thus approximating the real force field underlying the movement. The more a movement is curved, the more imprecisions related to this parametrization affect the model's performance. Finally, the original and extended models differed significantly in their predictions for all the error measures ($P < 0.001$).

We have further investigated whether our model captures the major temporal characteristics of the movement. We compared the VITE and F2REACH models' predictions to the real data for the peak amplitude, time at peak amplitude and speed asymmetry index, see Table 3. One way ANOVAs confirmed a very good match between our model's prediction and the data for all the above quantities (except for the velocity peak which was slightly lower, $P < 0.05$), whereas the predictions of the original VITE model differed significantly from the data ($P < 0.001$) except for the velocity peak amplitude. To illustrate the quality of the extended and original VITE models' predictions for the time-dependency of the signals, in Fig. 9 we compare instances of measured and predicted speed profiles (normalized in time). Finally we looked at the percentages of trajectory points comprised within the volumes defined by one and two standard deviations (SD) in order to evaluate the performance of the models at portions where the movement is very precise and systematic over trials (see Sect. 4.4 for details). The results show that 81% of the hand trajectory points predicted by our model were within two SDs of the mean trajectory against 40% of the points predicted by the original VITE model (Table 4 shows also the result for 1D).

One should emphasize that the F2REACH model generates these 3-dimensional movements using few parameters: β that controls the amplitude of the velocity and the two repulsive force vectors $\mathbf{v}$ and $\mathbf{w}$ (see Fig. 3b) that parameterize the force field surrounding the subject. The other two parameters α (time constant) and noise were fixed to 50 and 0.005 in all the simulations. The high accuracy with which the model manages to replicate the movements confirms that the model encapsulates the important features underlying free reaching movements. The force field is a key variable of the model. Next we show that the force field can be interpreted in relation to the bio-mechanical constraints of the subject's body.

5.3 Understanding the force field

Figure 10 shows the components of the virtual repulsive forces $\mathbf{v}$ and $\mathbf{w}$ parameterizing the force field of the F2REACH model (Eq. 2). We observe that the values of the components are clustered in two groups depending on the starting location of the movement. They are, thus, consistent

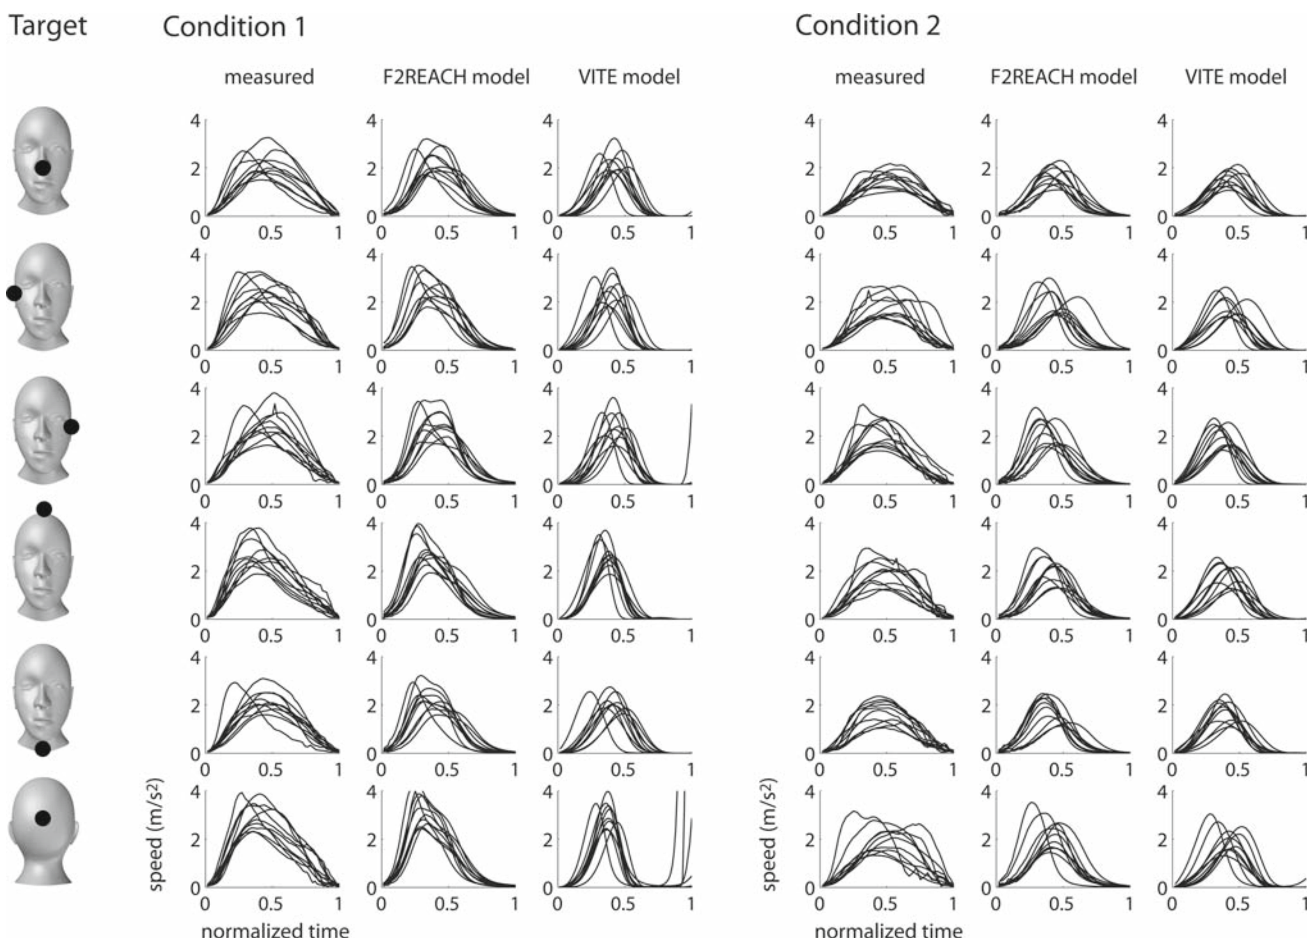

Fig. 9 Normalized time speed profiles of the measured human data and as predicted by the extended F2REACH and original VITE models, for the six target positions and the two initial conditions (see Fig. 5)

within the same condition (see Fig. 10a). The fact that movements to different targets are also clustered (forces underlying similar movements have similar components) suggests a certain regularity in the force field (see Fig. 10b). Finally, the trials related to one movement are clustered according to the subject executing the movement, which shows once more that the parameter values found for the repulsive force fields are not arbitrary (see Fig. 10c). Recall that the sign of the force vector governs the direction of the deviation and that, according to the expression of the modulating function h, the resulting force $\mathbf{F}(\mathbf{x}(t))$ coincides with $\mathbf{v}$ at the beginning of the movement and with $\mathbf{w}$ at the end of the movement, $\mathbf{F}(t = 0) = \mathbf{v}$ and $\mathbf{F}(\mathbf{x} = \mathbf{x}^*) = \mathbf{w}$.

Closer analysis of the clusters shows that the force $\mathbf{v}$, dominating the beginning of the movement, is highly dependent on the starting location in the x and y coordinates (see Fig. 4), whereas the force $\mathbf{w}$, dominating the end of the movement, varies according to the z direction. An intuitive explanation for this result is shown in Fig. 11 where we show the direction and amplitude of the repulsive forces $\mathbf{v}$, $\mathbf{w}$ and their modulated sum $\mathbf{F}(\mathbf{x}(t))$ (Eq. 3) for different types of movement. We see that, in the second starting position (arm extended in front of the body) the subject pushes his or her hand in the direction of the target (see Fig. 11a), whereas, in the first starting position (arm extended along the body) the subject must first push the hand to the right in order to avoid the body, and then bring the hand downwards in order to avoid reaching the limit of the shoulder joints (see Fig. 11b).

Table 4 Percentages of predicted trajectory points comprised within one and two standard deviation volumes (1 SD and 2 SD), see Fig. 6c, for the extended F2REACH and original VITE models

	F2REACH	VITE
1SD	54.64 ± 19.17	31.11 ± 8.44
2SD	80.89 ± 14.82	39.76 ± 9.18

This error measure is highly restrictive as it penalizes the model predictions at points where the five trials per subject and movement type are very consistent

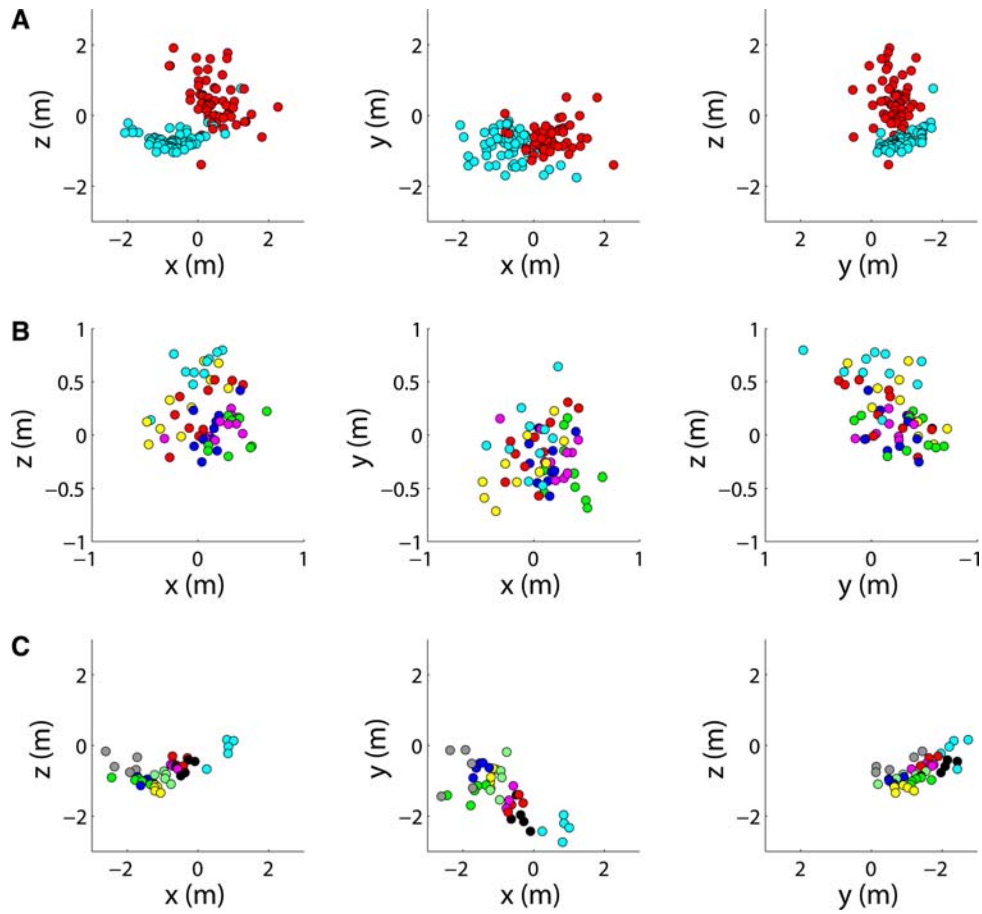

Fig. 10 Components of the repulsive forces **v** and **w**. **a** We show the components of the first repulsive force **v** for the two conditions: arm extended along the body (*red*) and arm extended in front of the body (*blue*). Two practically non-overlapping clusters can be observed showing a consistency of the parameter values within one condition. **b** We show the components of the second repulsive force **w** in the first condition for the six targets (different scale). Again the parameter values are clustered such that movements oriented toward one target are close together, showing a regularity in the repulsive force field. **c** We show the components of the first repulsive force **v** in the first condition and target right ear for the ten subjects. Clusters corresponding to the subjects can be identified for the five trials representing the movement

To better understand the effect of the forces when starting from the same initial condition, we compared the values found for the force components when reaching to two different targets (Fig. 11b, c). Unsurprisingly, the repulsive vector **v** is coherent across conditions irrespective of the target position, whereas the repulsive vector **w** depends on the target position and moves along the normal to the head surface at the target's position.

We also considered whether the magnitude of the repulsive force is related to the geometry of the subject's body, such as the length of the forearm for example. We observed a linear correlation between these two quantities (shown in Fig. 12): the shorter the arm, the more the hand must be pushed away to circumvent the head. Finally, we observed that the vectors of repulsive forces were coherent across subjects. These results are in agreement with the driving hypothesis of our model, namely that the curvature of reaching movements is the result of an explicit encapsulation of the task constraints in a control system which would, in the absence of constraints, produce straight-line motions. However, the opposite is not true, as we find non-null forces for quasi-straight movements, which are parallel to the motion. In the movements we have considered here, the task constraints comprise geometrical constraints related to the body.

5.4 Stability analysis of the model

The dynamical system described in Eq. 1 is globally asymptotically stable around a unique equilibrium point, the target position $\mathbf{x}^*$. We have omitted the analytical proof but the interested reader can convince themselves by computing the determinant of the Jacobian of the dynamical system around the fixed point and observe the latter to be always negative. Next we define the conditions under which the F2REACH model including the repulsive force field (see Eq. 2) is guaranteed to converge to the target. Let there be a perturbation that drifts the hand far away from the initial and target positions, such that $|\mathbf{x}(t) - \mathbf{x}(0)| - |\mathbf{x}(t) - \mathbf{x}^*(t)| < \varepsilon$, with $\varepsilon \in \mathbb{R}$ very small and $h(\mathbf{x}(t))$ and $1 - h(\mathbf{x}(t))$ approaching 1/2. The system converges to a stable state iff $|1/2(\mathbf{v} + \mathbf{w})| < 1$ such that the amplitude of the repulsive force field is smaller than the normalized attracting vector, i.e., the distance separating the target from the present position only gets smaller through time.[6] Note that all the forces' components found in our study satisfied the above condition.

[6] This result is, however, not valid in the vicinity of the initial position, which acts a second unstable attractor. Since this affects only a transient part of the motion (onset of the movement), which is unlikely to undergo perturbations, this could be ignored for the present study.

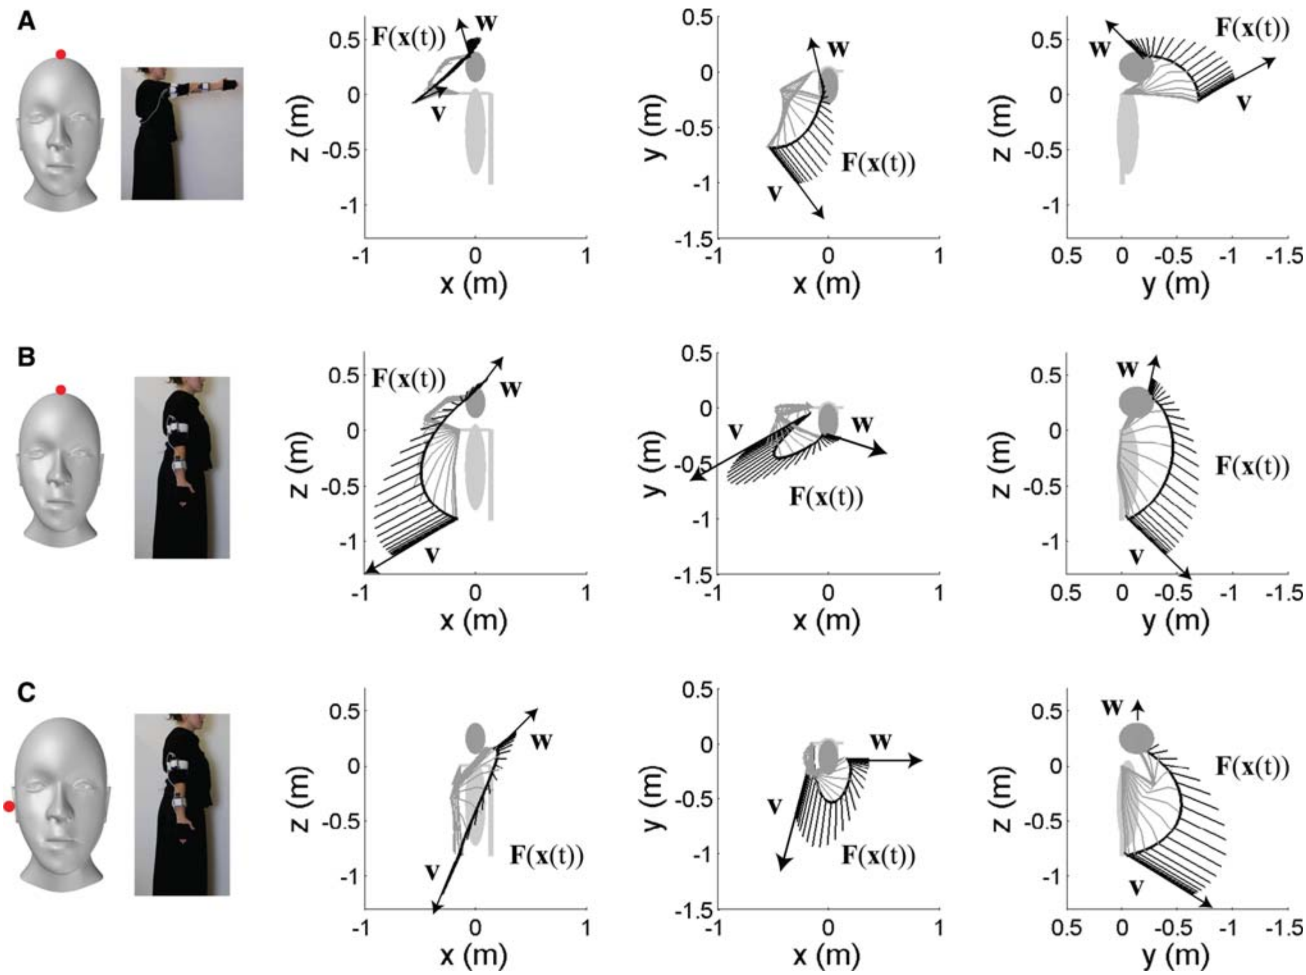

Fig. 11 Physical interpretation of the directions and amplitudes found for the repulsive forces **v**, **w** and their modulated sum $\mathbf{F}(\mathbf{x}(t))$ in our extended F2REACH model (Eqs. 2 and 3) relative to a schematized humanoid. Three movements of the same subject are shown. **a** The subject reaches for the top of the head (movement 4) with the arm extended in front of the body (condition 2). **b** Same position target as in **a** with the arm extended along the body (condition 1). **c** Same condition as in **b**, but the subject reaches to the left ear (movement 3). Due to the nature of the modulating function $h(\mathbf{x}(t))$, i.e., $h(\mathbf{x}(0)) = 1$ and $\lim_{t\to\infty} h(\mathbf{x}(t)) = 0$ (see the Appendix), the resulting force $\mathbf{F}(\mathbf{x}(t))$ coincides with **v** at the beginning of the movement and with **w** at the end of the movement, i.e., $\mathbf{F}(t = 0) = \mathbf{v}$ and $\mathbf{F}(\mathbf{x} = \mathbf{x}^*) = \mathbf{w}$ (Eq. 3). From **a** and **b** one can see that the initial condition affects the repulsive forces **v** and **w**. For example, in the second condition (**a**), **v** is in the direction of the target, whereas in the first condition (**b** and **c**) it is deviated to the right in order to avoid the body and downwards such that the arm does not reach the shoulder extension limit. In addition, **v** is coherent within the same condition (see **b** and **c**). The target position particularly affects the repulsive force **w** (predominant at the end of the movement) that is similar to the normal of the head surface approached. $\mathbf{F}(\mathbf{x}(t))$ was scaled for illustrative reasons

6 Discussion

We have hypothesized that the curvature of unconstrained reaching movements is due to an explicit encapsulation of the task constraints by the CNS in a virtual force field (F2). Movements thus unfold in time according to a dynamical system that attracts the hand to the target position while repelling it from undesirable locations in space (such as objects in the environment, the subject's body and joint limits) and while compensating for unexpected perturbations of the arm. Furthermore, we have argued that the curvature observed in natural movements is not a by-effect of the inherent dynamics of the body but a necessary and voluntarily controlled feature.

In order to probe our hypothesis, we have conducted motion studies in which healthy adult subjects produced natural reaching motions directed to various locations on their head. To highlight the effect that body constraints may have on the curvature of the movement, we asked the subjects to initiate the movement from two locations: one that required the subject to move alongside the body, the other which allowed the subject to move quasi freely. We showed that our mathematical model, the F2REACH model, could predict the major kinematic features of the movements, such as the bell-shaped velocity profile. Most importantly, it could account for both the weak and strong curvatures of the movements.

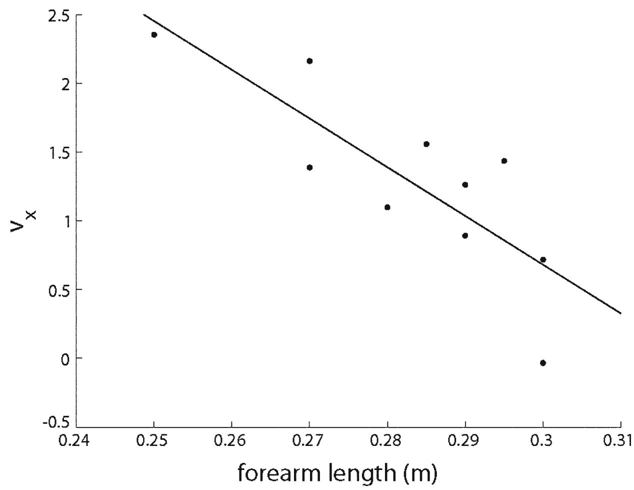

Fig. 12 The amplitude of the repulsive force **w** (Eqs. 2 and 3) is linearly correlated to the length of the subjects' forearms for the movement reaching to the left ear with the right arm (condition 1), showing that the repulsive forces in our model are affected by geometrical features of the body. Intuitively, with a shorter forearm, the hand needs to be pushed stronger away to circumvent the head

This led us to argue that a single controller underlies both straight and curved movements. The controller adapts the trajectory according to multiple constraints the subject has consciously or not decided to take into account. Although we have only shown that it can precisely reproduce the kinematics of self-oriented movements, the model is general and can generate natural movements to any target in both intrapersonal and extrapersonal spaces, e.g., in another study that investigates imitation of unnatural postures we successfully use this model to simulate reaching to objects on a table.

6.1 Assumptions of the model

The first assumption we have followed is that of a *functionally hierarchical motor control system* proposed by Bernstein (1947) and translated partly in Bernstein (1996). The hierarchy consists of four levels: complex actions with abstract goals, dealing with 3-dimensional space, muscle synergies, posture and muscle tone. In our study, we considered the first and second levels, in that we addressed 3-dimensional goal-directed reaching movements, characterized by a single target position. By leaving out the question of how such high-level control is then translated into muscle synergies and the control of posture and muscle tone, we follow the observation that: electrical stimulation of the brain motor area elicits reaching movements in primates (Graziano et al. 2002, 2005) and leg movements in frogs (Bizzi et al. 1982). Interestingly, all of these movements converge to the same position in extrinsic space independently from the initial posture. Thus, the control of these movements seems to use solely the definition of the desired final position, and not a description of low-level muscle activations (in a way functionally similar to muscle synergies when compared to activating individual muscles, see d'Avella et al. (2003)). In addition these studies indicate that reaching movements are extensively represented in the motor cortex.

Another argument in favor of a "high-level" extrinsic 3-dimensional representation of movements come from evidence of the many to one mappings between: (1) muscles and joint configurations, (2) muscles and end-effector positions or (3) joint configurations and end-effector positions. Controlling the hand in a 3-dimensional extrinsic space over an intrinsic joint space is advantageous in that it allows to easily encapsulate task constraints, such as avoiding surrounding objects, and plan movements accordingly (these task constraints would have an infinite number of possible representations in the joint and muscle spaces). Also note that we have assumed that movements were computed in a Cartesian frame of reference located on the body. It would however be conceivable to compute the same movement according to a polar coordinate system without affecting the prediction of the model.

The fact that we do not address the above two lower-levels of motor control, is a limitation of the model. As stated by Bernstein, the problem of translating a kinematic signal encoded in a 3-dimensional extrinsic frame of reference into muscle activations (so-called degrees of freedom problem) is complex because of the redundancy of the muscular system. An infinity of different muscle activations leads to the same kinematic motion of the end-effector. Although this problem is of the highest importance for a complete motor control theory, we do not address this problem here [see d'Avella et al. (2003), Todorov and Jordan (2002) and Guigon et al. (2007) for possible solutions].

Another important assumption we make is that the CNS can represent forces internally. Our model is based upon a force field that encapsulates the constraints of the task, which implies the knowledge of a mapping between different locations in the subject's peripersonal space and virtual repulsive forces. It thus requires the existence of an internal model of the environment in terms of attractive or repulsive force fields in the brain. The above hypotheses are not at odds with the literature. There is substantial evidence that the brain is capable of learning an internal representation of external forces in order to adapt its control of the motion of the hand (Shadmehr and Mussa-Ivaldi 1994; Conditt et al. 1997; Shadmehr and Brashers-Krug 1997; Thoroughman and Shadmehr 2000; Gandolfo et al. 1996), when subjected to these for a long enough period of time. Another force that is centrally represented and integrated in the internal dynamic control models for reaching is the gravitational force (Shadmehr and Mussa-Ivaldi 1994; Papaxanthis et al. 1998). In our model the geometry of the body and external objects, among other factors, contribute to the force field.

Accordingly, McIntyre et al. (1995) have shown that the brain may integrate an external constraint such as a curved surface through an a priori internal model of the surface geometry.

6.2 Properties of the model

Interesting properties of the F2REACH model for motor control are: (i) the system is *asymptotically and globally stable*; (ii) it exploits a biologically plausible *signal-dependant noise* and (iii) planning of the movement is done through closed-loop control. This enables on-the-go re-computation of the motion in the face of perturbation or imprecision in the sensory-motor loop. Closed-loop control through afferent and efferent internal feedback loops (Desmurget and Grafton 2000) allows to take into account the uncertainty of the "real-world" instead of just "playing a prerecorded tape" (Todorov 2004). We suggest that only an online mechanism that tightly couples movement planning with movement execution could explain the irregular curvatures observed in some of the trials; the latter were likely due to an on-the-go correction of the trajectory.

Most importantly, we have proposed a force field framework as a powerful mechanism for integrating various constraints related to, e.g., the dynamics and geometry of the arm, external objects and the person's own embodiment, into a unique and generic controller. Whereas the goal of the controller is encoded according to kinematic variables (a position to reach), the constraints are encoded in dynamic variables, the force field, and may as well be expressed in an intrinsic (limit joint angles) or extrinsic (surrounding objects) frame of reference. This framework could reconcile findings that argue for both dynamic and kinematic planning (Vetter et al. 2002; Admiraal et al. 2004), in providing a computational account for how the dynamics of the arm can be taken into account in kinematic planning (Sabes and Jordan 1997). It also explains how external objects might influence the trajectory of the hand (Brenner and Smeets 1995).

Furthermore, the representation of this environmental force field generalizes to performing the motion faster or slower (Harris and Wolpert 1998). This is equivalent to learning to vary the value of the factor β in our model (see Eq. 1). Finally the representation of the force field, although local, extends to nearby locations (smoothly decaying away from the position of the perturbation). Similarly, our expression of the force field is spatially continuous.

The extent to which the model's predictions can be generalized to any reaching movement remains to be shown, since we only demonstrated a good agreement of the model with data from reaching movements directed to the head. The movements we have addressed are nevertheless quite generic in that they were entirely unconstrained. For example, we did not observe a reduction of the degrees of freedom as in Klein Breteler et al. (1998) where the subjects had a tendency to produce movements in 2D rather than in 3D (see example in Fig. 1). In addition many of the velocity profiles recorded, exhibited asymmetric velocity profiles similar to those observed (Gielen et al. 1985; Brown and Cooke 1990). These characteristics are present in all reaching movements and we are thus confident that the model is generic in its representation of the class of reaching movements.[7]

The force field in our model is parameterized by two constant forces and is thus only an approximation of the real underlying force field. This approximation may lead to imprecisions in the model's predictions, especially in places where the field changes importantly locally.

While our model proposes a way in which the brain may encapsulate all types of motion-related constraints (e.g., body and joint-limits avoidance, inertia of the arm) within a general controller of reaching movements, we do not provide a general method for expressing these constraints in the form of a force field. Our future efforts will concentrate on segmenting the contributions of different constraints and on a mechanism that would allow to learn these through experience.

6.3 Predictions of the model

Our model is consistent with several experimental observations and provides a theoretical basis for their interpretation.

For example, in different pointing and reaching studies, systematic misdirections of the fingertip trajectory were observed (de Graaf et al. 1991, 1994; Brenner and Smeets 1995). The misdirections were clockwise and anticlockwise when pointing to targets on the right and on the left frontal space, respectively. To explain their results the authors hypothesized a distorted and contracted internal representation of space (de Graaf et al. 1991, 1994) or speculated that the subjects anticipate the purpose of the target (Brenner and Smeets 1995). Within the repulsive force field framework we propose in this paper, these misdirections are created by the geometrical relationship between the target and the subject's body.[8] Our model predicts that if one was to repeat the experiment in a different part of the workspace where the misdirections are mainly due to body avoidance, the misdirections would be anticlockwise and clockwise when the target is respectively

[7] Current work of ours has applied the model to account for reaching movements oriented to targets on a table in natural and unnatural postures where an artificial constraint is introduced. Preliminary results show that the model again encapsulates with high accuracy all the features of the movements (unpublished data).

[8] These two similar studies, de Graaf et al. (1991) and Brenner and Smeets (1995), puzzlingly reported different results. We suggest that the differences observed can be attributed to the distance chosen from the subject to the initial position of the hand [25 cm in Brenner and Smeets (1995) and 40 cm in de Graaf et al. (1991)], as the repulsive forces responsible for avoiding the body would fade away as this distance increases.

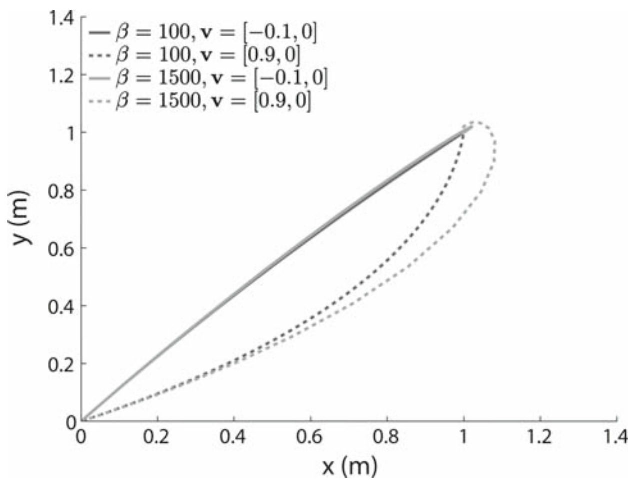

Fig. 13 Prediction of the F2REACH model: the curvature increases with higher speed, here equivalent to higher β values. The effect is not visible in quasi-straight movements

right and left from the closest virtual line connecting the subject's trunk with the hand's initial position.

Furthermore, our model predicts that faster movements may be more curved, as shown on Fig. 13. Even though this prediction has been empirically observed (Klein Breteler et al. 1998), it contradicts several experimental and theoretical studies that have shown curvature-speed invariance (Nishikawa et al. 1999; Sha et al. 2006; Liebermann et al. 2008) and suggest that speed and path are planned independently (Todorov and Jordan 2002; Torres and Zipser 2002; Biess et al. 2007). In our model speed modulates the curvature of the path by construction. However, the deviation is also proportional to the magnitude of the repulsive force field such that this effect is particularly important for highly curved movements (see Fig. 13). This might explain why curvature-speed invariance is more consistently observed, as highly curved movements are rarely studied. Otherwise, an additional compensatory mechanism should be added to the model that modulates the force field as a function of desired speed.

Finally, the model suggests that the asymmetry of the velocity profile is due to the difference in directions between the repulsive force field and attracting vector. Finally, even though the curvature of a movement is highly systematic and reproducible (Soechting and Lacquaniti 1981; Pellegrini and Flanders 1996; Admiraal et al. 2004), our model would predict that if you alter the geometry of the subject's body, such as adding a false belly for example, then reaching movements will be displaced away from the artificial object even if this object does not interfere with the original trajectory. Our model also predicts that the shape of the object would matter.

6.4 Neural correlates of the model

Most importantly, the F2REACH model we propose is compatible with neurophysiological studies. Primate brain areas have been identified as the loci of the computations involved in the original VITE model (Bullock et al. 1998). Specifically, it was shown that the model's variables display the same dynamics of activation (e.g., response profiles and latency of activity onset) as that of populations of neurons: the hand velocity might be represented in area 4, whereas the hand acceleration and position in area 5. Note that the extended F2REACH model solicits only quantities that would be easily accessible to the CNS such as distances from the target and initial positions.

A novel feature of the model is the repulsive force field that shapes the landscape of the workspace, meaning that not each position is equally likely to be visited. In other words, the model assumes the existence of neural populations coding for forces related to the body and surrounding objects. Area 5 is a putative region for the computation of the force field, as it receives abundant somatosensory and visual inputs that are necessary for the encapsulation of the geometrical properties of the body and surrounding objects in an internal model (Scott et al. 1997; Graziano et al. 2000). We thus predict the existence of a population of neurons in area 5, whose activity would be close to baseline during straight movements and would rotate in curved movements. In addition, the activation of these neurons would be modulated by the introduction of new objects in the workspace.

6.5 Conclusion

We showed that not only the spatial, but also the temporal features of unconstrained and naturally curved reaching movements could be modeled through a dynamical system modulated by a virtual force-field. We found that the model was in very good agreement with kinematic data from human motions, during unconstrained reaching movements directed to the head. We showed that the natural curvature of these movements could be attributed to the interplay between a target attractor and virtual repulsive forces that encapsulate a representation of the geometry of the subject's body. Such a representation is a simple and powerful way to generate kinematically-driven trajectories that comply with the underlying dynamic constraints.

7 Appendix

7.1 Original VITE system

The original VITE model's dynamics as given by Bullock and Grossberg (1988):

$$\dot{\mathbf{y}}(t) = \alpha(-\mathbf{y}(t) + \mathbf{x}^*(t) - \mathbf{x}(t))$$
$$\dot{\mathbf{x}}(t) = \beta t^{\nu} \mathbf{y}(t) \tag{4}$$

where $\mathbf{x}(t)$ corresponds to the current position of the hand in a three-dimensional extrinsic frame of reference and $\mathbf{x}^*(t)$ is the location of the target (see Fig. 4). $\mathbf{y}$ is a secondary variable related to the hand velocity. α and β are real positive time constants and ν is a real positive exponent parameter. The model recomputes at each time step the hand position $\mathbf{x}(t)$, so as to generate an overall straight trajectory to the target that follows a bell-shaped velocity profile. The first term of the equation ensures that the unprimed acceleration vector $\dot{\mathbf{y}}(\mathbf{t})$ is always directed toward the target, i.e., $\mathbf{x}^*(\mathbf{t}) - \mathbf{x}(t)$, so that the target's position $\mathbf{x}^*$ forms a unique attractor of the system. The amplitude of the acceleration $\dot{\mathbf{y}}(\mathbf{t})$ is proportional to the distance separating the hand and the target. $\mathbf{y}(t)$ grows quickly at the beginning of the movement and slows down exponentially towards the end of the movement. To compensate for this asymmetric velocity profile, $\mathbf{y}(t)$ is scaled down in the second equation by a time-dependent variable βt^{ν}. $\dot{\mathbf{x}}(\mathbf{t})$ is the hand's velocity and can be viewed as the output activity of a corresponding neural population that would control agonist muscle motoneurons (Bullock and Grossberg 1988).

7.2 Nonlinear functions used in the F2REACH model

The form of the nonlinear function $\mathbf{g}$ in Eq. 1 is the following:

$$\mathbf{g}(\mathbf{u}) = |\mathbf{d} - \mathbf{u}|\mathbf{u} \tag{5}$$

where the control vector $\mathbf{u}(\mathbf{t}) = \mathbf{x}^*(t) - \mathbf{x}(t) + \boldsymbol{\eta}$ is the vector separating the actual hand position $\mathbf{x}(t)$ from the desired hand position $\mathbf{x}^*(t)$ (does not need to be stationary) with signal dependant noise $\boldsymbol{\eta}$. The operator $|\ |$ stands for the norm of the vector and $\mathbf{d}$ is defined as:

$$\mathbf{d}(t) = \mathbf{x}^*(t) - \mathbf{x}(0) \tag{6}$$

the vector between the target $\mathbf{x}^*(t)$ and initial position $\mathbf{x}(0)$, such that the term $|\mathbf{d} - \mathbf{u}|$ is equivalent to the distance separating the actual position of the end-effector from its initial position. t is set to 0 each time a new movement is initiated. In the absence of noise in the control signal $\mathbf{u}$, the multiplicative factor $|\mathbf{d} - \mathbf{u}|$ would be 0 at $t = 0$ and no movement would be initiated.

The function h that modulates the force field in Eq. 3 is defined by:

$$h(\mathbf{u}) = \frac{|\mathbf{u}|}{|\mathbf{u}| + |\mathbf{d} - \mathbf{u}|} \tag{7}$$

and normalizes the amplitude of the control signal $\mathbf{u}$.

Acknowledgments This work was funded by the by the Sport and Rehabilitation Engineering Program at EPFL and EU Project IST-2004-004370 RobotCub.

References

Admiraal MA, Kusters MJMAM, Gielen SCAM (2004) Modeling kinematics and dynamics of human arm movements. Motor Control 8(3):312–338
Ajemian R, Bullock D, Grossberg S (2001) A model of movement coordinates in the motor cortex: posture-dependent changes in the gain and direction of single cell tuning curves. Cereb Cortex 11(12):1124–1135
Atkeson CG, Hollerbach JM (1985) Kinematic features of unrestrained vertical arm movements. J Neurosci 5(9):2318–2330
Bernstein NA (1947) On the construction of movements. Medgiz, Moscow
Bernstein NA (1996) Levels of construction of movements. In: Latash ML, Turvey MT (eds) Dexterity and its development. Lawrence Erlbaum Associates, Mahwah, NJ, pp 115–170
Biess A, Liebermann DG, Flash T (2007) A computational model for redundant human three dimensional pointing movements: integration of independant spatial and temporal motor plans simplifies movement dynamics. J Neurosci 27(48):13045–13064
Bizzi E, Accornero N, Chapple W, Hogan N (1982) Arm trajectory formation in monkeys. Exp Brain Res 46(1):139–143
Brenner E, Smeets JBJ (1995) Moving one's finger to a visually specified position: target orientation influences the finger's path. Exp Brain Res 105(2):318–320
Brown SH, Cooke JD (1990) Movement-related phasic muscle activation. I. Relations with temporal profile of movement. J Neurophysiol 63(3):455–464
Bullock D, Grossberg S (1988) Neural dynamics of planned arm movements: emergent invariants and speed-accuracy properties during trajectory formation. Psychol Rev 95(1):49–90
Bullock D, Grossberg S, Mannes C (1993) A neural network model for cursive script production. Biol Cybern 70(1):15–28
Bullock D, Cisek P, Grossberg S (1998) Cortical networks for control of voluntary arm movements under variable force conditions. Cereb Cortex 8(1):48–62
Christel MI, Billard A (2002) Comparison between macaques' and humans' kinematics of prehension: the role of morphological differences and control mechanisms. Behav Brain Res 131(1–2):169–184
Cisek P, Grossberg S, Bullock D (1998) A cortico-spinal model of reaching and proprioception under multiple task constraints. J Cogn Neurosci 10(4):425–444
Clamman HP (1969) Statistical analysis of motor unit firing pattern in human skeletal muscle. Biophys J 9:1233–1251
Clancy EA, Hogan N (1999) Probability density of the surface electromyogram and its relation to amplitude detectors. IEEE Trans Biomed Eng 46(6):730–739
Conditt MA, Gandolfo F, Mussa-Ivaldi FA (1997) The motor system does not learn the dynamics of the arm by rote memorization of past experience. J Neurophysiol 78(1):554–560
Cruse H, Brüwer M (1987) The human arm as a redundant manipulator: the control of path and joint angles. Biol Cybern 57(1–2):134–144
d'Avella A, Saltiel P, Bizzi E (2003) Combinations of muscle synergies in the construction of a natural motor behavior. Nat Neurosci 6(3):300–308
de Graaf JB, Sittig AC, Denier von der Gon JJ (1991) Misdirections in slow goal-directed arm movements and pointer-setting tasks. Exp Brain Res 84(2):434–438
de Graaf JB, Sittig AC, Denier von der Gon JJ (1994) Misdirections in slow, goal-directed arm movements are not primarily visually based. Exp Brain Res 99(3):464–472
Dalby JT, Gibson D, Grossi V, Schneider RD (1980) Lateralized hand gesture during speech. J Motor Behav 12(4):292–297
De Renzi E, Lucchelli F (1988) Ideational apraxia. Brain 111(Pt 5):1173–1185

Desmurget M, Grafton S (2000) Forward modeling allows feedback control for fast reaching movements. Trends Cogn Sci 4(11):423–431

Desmurget M, Jordan M, Prablanc C, Jeannerod M (1997) Constrained and unconstrained movements involve different control strategies. J Neurophysiol 77(3):1644–1650

Fitts PM (1954) The information capacity of the human motor system in controlling the amplitude of movement. J Exp Psychol 47(6):381–391

Flash T (1987) The control of hand equilibrium trajectories in multi-joint arm movements. Biol Cybern 57(4–5):257–274

Flash T, Hogan N (1985) The coordination of arm movements: an experimentally confirmed mathematical model. J Neurosci 5(7):1688–1703

Gandolfo F, Mussa-Ivaldi FA, Bizzi E (1996) Motor learning by field approximation. Proc Natl Acad Sci USA 93(9):3843–3846

Gielen CCAM, van den Oosten K, Pull ter Gunne F (1985) Relation between EMG activation patterns and kinematic properties of aimed arm movements. J Mot Behav 17(4):421–442

Goldenberg G, Hagmann S (1997) The meaning of meaningless gestures: a study of visuo-imitative apraxia. Neuropsychologia 35(3):333–341

Graziano MSA, Cooke DF, Taylor CSR (2000) Coding the location of the arm by sight. Science 290(5497):1782–1786

Graziano MSA, Taylor CSR, Moore T (2002) Complex movements evoked by microstimulation of precentral cortex. Neuron 34(5):841–851

Graziano MSA, Aflalo TNS, Cooke DF (2005) Arm movements evoked by electrical stimulation in the motor cortex of monkeys. J Neurophysiol 94(6):4209–4223

Guigon E, Baraduc P, Desmurget M (2007) Computational motor control: redundancy and invariance. J Neurophysiol 97(1):331–347

Hamilton AFC, Wolpert DM (2002) Controlling the statistics of action: obstacle avoidance. J Neurophysiol 87(5):2434–2440

Harris CM, Wolpert DM (1998) Signal-dependent noise determines motor planning. Nature 394(6695):780–784

Hersch M, Billard A (2007) Reaching with multi-referential dynamical systems. Auton Robot 25(1–2):71–83

Hoos H, Stützle T (2004) Stochastic local search: Foundations and applications. Elsevier/Morgan Kaufmann, San Francisco

Klein Breteler MD, Meulenbroek RGJ, Gielen SCAM (1998) Geometric features of workspace and joint-space paths of 3D reaching movements. Acta Psychol 100(1–2):37–53

Lacquaniti F, Terzuolo C, Viviani P (1983) The law relating the kinematic and figural aspects of drawing movements. Acta Psychol 54(1–3):115–130

Lacquaniti F, Soechting JF, Terzuolo SA (1986) Path constraints on point-to-point arm movements in three-dimensional space. Neuroscience 17(2):313–324

Lavergne J, Kimura D (1987) Hand movement asymmetry during speech: no effect of speaking topic. Neuropsychologia 25(4): 689–693

Liebermann DG, Biess A, Gielen CCAM, Flash T (2006) Intrinsic joint kinematic planning. II: Hand-path predictions based on a Listing's plane constraint. Exp Brain Res 171(2):155–173

Liebermann DG, Krasovsky T, Berman S (2008) Planning maximally smooth hand movements constrained to nonplanar workspaces. J Mot Behav 40(6):516–531

Matthews PBC (1996) Relationship of firing intervals of human motor units to the trajectory of post-spike after-hyperpolarization and synaptic noise. J Physiol 492(Pt 2):597–628

McIntyre J, Gurfinkel EV, Lipshits MI, Droulez J, Gurfinkel VS (1995) Measurements of human force control during a constrained arm motion using a force-actuated joystick. J Neurophysiol 73(3):1201–1222

Miall RC, Haggard PN (1995) The curvature of human arm movements in the absence of visual experience. Exp Brain Res 103(3):421–428

Moran DW, Schwartz AB (1999) Motor cortical activity during drawing movements: population representation during spiral tracing. J Neurophysiol 82(5):2693–2704

Morasso P (1981) Spatial control of arm movements. Exp Brain Res 42(2):223–227

Nakano E, Imamizu H, Osu R, Uno Y, Gomi H, Yoshioka T, Kawato M (1999) Quantitative examinations of internal representations for arm trajectory planning: minimum commanded torque change model. J Neurophysiol 81(5):2140–2155

Neter J, Kutner MH, Nachtsheim CJ, Wasserman W (1996) Applied linear statistical models. McGraw-Hill/Irwin, Chicago

Nishikawa KC, Murray ST, Flanders M (1999) Do arm postures vary with the speed of reaching? J Neurophysiol 81(5):2582–2586

Okadome T, Honda M (1999) Kinematic construction of the trajectory of sequential arm movements. Biol Cybern 80(3):157–169

Oldfield RC (1971) The assessment and analysis of handedness: the Edinburgh inventory. Neuropsychologia 9(1):97–113

Osu R, Uno Y, Koike Y, Kawato M (1997) Possible explanations for trajectory curvature in multijoint arm movements. J Exp Psychol Hum Percept Perform 23(3):890–913

Osu R, Kamimura N, Iwasaki H, Nakano E, Harris CM, Wada Y, Kawato M (2004) Optimal impedance control for task achievement in the presence of signal-dependant noise. J Neurophysiol 92(2):1199–1215

Paine RW, Grossberg S, Van Gemmert AWA (2004) A quantitative evaluation of the AVITEWRITE model of handwriting learning. Hum Mov Sci 23(6):837–860

Papaxanthis C, Pozzo T, Popov KE, McIntyre J (1998) Hand trajectories of vertical arm movements in one-G and zero-G environments. Evidence for a central representation of gravitational force. Exp Brain Res 120(4):496–502

Pellegrini JJ, Flanders M (1996) Force path curvature and conserved features of muscle activation. Exp Brain Res 110(1):80–90

Petreska B, Adriani M, Blanke O, Billard AG (2007) Apraxia: a review. Prog Brain Res 164:61–83

Sabes PN, Jordan MI (1997) Obstacle avoidance and a perturbation sensitivity model of motor planning. J Neurosci 17(18):7119–7128

Sauser E, Billard A (2006) Dynamic updating of distributed neural representations using forward models. Biol Cybern 95(6):567–588

Schaal S, Sternad D (2001) Origins and violations of the 2/3 power law in rhythmic three-dimensional arm movements. Exp Brain Res 136(1):60–72

Schmidt RA, Zelaznik H, Hawkins B, Frank JS, Quinn JT (1979) Motor-output variability: a theory for the accuracy of rapid motor acts. Psychol Rev 47(5):415–451

Scott SH, Sergio LE, Kalaska JF (1997) Reaching movements with similar hand paths but different arm orientations. II. Activity of individual cells in dorsal premotor cortex and parietal area 5. J Neurophysiol 78(5):2413–2426

Sha D, Patton J, Mussa-Ivaldi FA (2006) Minimum jerk reaching movements of human arm with mechanical constraints at endpoint. Int J Comput Syst Signal 7(1):41–50

Shadmehr R, Brashers-Krug T (1997) Functional stages in the formation of human long-term motor memory. J Neurosci 17(1): 409–419

Shadmehr R, Mussa-Ivaldi FA (1994) Adaptive representation of dynamics during learning of a motor task. J Neurosci 14(5 Pt 2): 3208–3224

Soechting JF, Lacquaniti F (1981) Invariant characteristics of a pointing movement in man. J Neurosci 1(7):710–720

Soechting JF, Buneo CA, Herrmann U, Flandres M (1995) Moving effortlessly in three dimensions: does Donders' law apply to arm movement? J Neurosci 15(9):6271–6280

St-Amant Y, Rancourt D, Clancy EA (1998) Influence of smoothing window length on electromyogram amplitude estimates. IEEE Trans Biomed Eng 45(6):795–800

Sutton GG, Sykes K (1967) The variation of hand tremor with force in healthy subjects. J Physiol 191(3):699–711

Thoroughman KA, Shadmehr R (2000) Learning of action through adaptive combination of motor primitives. Nature 407(6805):742–747

Todorov E (2004) Optimality principles in sensorimotor control. Nat Neurosci 7(9):907–915

Todorov E, Jordan MI (2002) Optimal feedback control as a theory of motor coordination. Nat Neurosci 5(11):1226–1235

Torres EB, Zipser D (2002) Reaching to grasp with a multi-joined arm. I. Computational model. J Neurophysiol 88(5):2355–2367

Uno Y, Kawato M, Suzuki R (1989) Formation and control of optimal trajectory in human multijoint arm movement. Minimum torque-change model. Biol Cybern 61(2):89–101

Vetter P, Flash T, Wolpert DM (2002) Planning movements in a simple redundant task. Curr Biol 12(6):488–491

Wada Y, Kaneko Y, Nakano E, Osu R, Kawato M (2001) Quantitative examinations for multi joint arm trajectory planning—using a robust calculation algorithm of the minimum commanded torque change trajectory. Neural Netw 14(4–5):381–393

Wang W, Chan SS, Heldman DA, Moran DW (2007) Motor cortical representation of position and velocity during reaching. J Neurophysiol 97(6):4258–4270

Wolpert DM, Ghahramani Z, Jordan MI (1994) Perceptual distortion contributes to the curvature of human reaching movements. Exp Brain Res 98(1):153–156

Wolpert DM, Ghahramani Z, Jordan MI (1995) Are arm trajectories planned in kinematic or dynamic coordinates? An adaptation study. Exp Brain Res 103(3):460–470

Dynamical System Modulation for Robot Learning via Kinesthetic Demonstrations

Micha Hersch, Florent Guenter, Sylvain Calinon, and Aude Billard
Learning Algorithms and Systems Laboratory - LASA
School of Engineering - EPFL
Station 9 - 1015 Lausanne - Switzerland

Abstract

We present a system for robust robot skill acquisition from kinesthetic demonstrations. This system allows a robot to learn a simple goal-directed gesture, and correctly reproduce it despite changes in the initial conditions, and perturbations in the environment. It combines a dynamical system control approach with tools of statistical learning theory and provides a solution to the inverse kinematics problem, when dealing with a redundant manipulator. The system is validated on two experiments involving a humanoid robot: putting an object into a box, and reaching for and grasping an object.

Index Terms

Robot Programming by Demonstration, Dynamical System Control, Gaussian Mixture Regression

Corresponding author: Micha Hersch, micha.hersch@epfl.ch
Conditionally accepted short paper submitted to the IEEE Transactions on Robotics
Paper # A07-471/A06-417

This work was supported by the European Commission through the EU Integrated Projects ROBOT-CUB (FP6-IST-004370), Cogniron (FP6-IST-FET-002020) and ROBOT@CWE (FP6-034002).

Dynamical System Modulation for Robot Learning via Kinesthetic Demonstrations

I. Introduction

As robots are progressively coming out of the controlled environment of assembly lines to pervade the much less predictable domestic environments, there is a need to develop new kinds of controllers that can cope with changing environments and that can be taught by unskilled human users. In order to address this last issue, *Programming by Demonstration* (PbD) has emerged as a promising approach [1]. PbD has been mostly used in two cases: for tasks involving no or very loose interaction with the environment (like writing, martial arts or communicative gestures) human demonstrations are used to train a movement model, which can be used to reproduce the task. Those movement models (also used in computer animation or visual gesture recognition) usually imply some averaging process (LWR [2], HSTMM [3]), possibly in a latent space (GPLVM [4], ST-Isomap [5]) or some probabilistic model like Bayesian Networks [6]. And for more complex tasks, involving precise interactions with the environment, the robot learns from examples how to sequence a set of hard-coded controllers for a given task. This has been done using HMMs [7] or knowledge-based systems [8].
In our work, we position ourselves in between those two approaches and combine learning of a task-dependent modulation of a built-in controller. We, thus, start with a basic built-in controller (or motion primitive) that consists in a dynamical system with a single stable attractor. We modulate the trajectories generated by this controller to be task-specific, by learning a probabilistic model of the task-based trajectory, as shown by a human user. This results in a general framework for learning and reproducing goal-directed gestures, despite different initial conditions and changes occurring during task execution. In this respect, we improve in several ways classical control approaches for goal-directed motions, such as [9].

The closest work to ours is [2], which uses a dynamical system for goal-directed reaching. We depart from this work in two ways: First, we propose a hybrid controller composed of two of our basic dynamical systems working concurrently in end-effector and joint angle spaces. This results in a controller that has no singularities. Second, the dynamical system approach gives us a controller robust in the face of perturbations, which can recompute the trajectory on-line to adapt to sudden displacements of the target or unexpected motion of the arm during motion, and we provide experimental results on the robustness to static and dynamic changes in the environment. While our controller is less precise than ad-hoc controller (e.g. [10]), it is more general in that it can be easily modulated to achieve arbitrary goal-directed reaching tasks.

In the experiments presented here, the motions are demonstrated to the robot by a human user moving the robots' limbs passively (kinesthetic training). In Section IV, we validate the approach on two different tasks, namely placing an object into a box, and reaching-to-grasp a chess piece, see Fig. 2 for illustrations of these two tasks.

II. Overview

The system is designed to enable a robot to learn to modulate its generic controller to produce any arbitrary goal-directed motion. The model must be generic so as to reproduce the motion given different initial conditions and under perturbations during execution. Moreover, the architecture of the system must permit the use of different control variables for encoding the motion. Here, we compare a control in either velocity or acceleration. We refer to those further as the *velocity model* (see Section II-B) and the *acceleration model* (see Section II-C).

A. System Architecture

The structure of the system is the same for both models and is schematized in Fig. 1. During training, the relevant variables (end-effector velocity profiles for the velocity model, or end-effector positions, velocities and accelerations for the acceleration model) are extracted from the set of demonstrated trajectories and used to train a Gaussian Mixture Model (GMM) (see Table I). During reproduction, the trajectory is specified by a spring-and-damper dynamical system modulated by the GMM (see section III). The target is tracked by a stereo-vision system and is set to be the attractor point of the dynamical system. At each time step, the desired velocity computed by the model is then fed to a PID controller for execution. This does not hinder the online adaptation of the movement.

B. Velocity Model

The first way to encode a motion in a GMM, is to consider the velocity profile of the end-effector as a function of time $\dot{\mathbf{x}}(\mathbf{t})$. Thus, the input variable ζ is the time and the output variable ξ is the velocity, like in the following velocity model:

$$\dot{\mathbf{x}}^{\mathbf{m}} = \tilde{\mathcal{F}}_{\dot{\mathbf{x}}}(t) \tag{2}$$

In other words, the movement is modeled as a velocity profile, given by a function of time, which is learned as described in Table I. Here and henceforth, $\dot{\mathbf{x}}^{\mathbf{m}} \in \mathbb{R}^m$ is the end-effector velocity specified by the task model. $\tilde{\mathcal{F}}_{\dot{\mathbf{x}}}$ is obtained by applying (1) with the appropriate variables.

TABLE I
SUMMARY OF GAUSSIAN MIXTURE REGRESSION (GMR).

GMR is a method suggested by [11] for statistically estimating a function $\mathcal{F}_\xi(\zeta)$ given by a "training set" of N examples $\{(\zeta^i, \xi^i)\}_{i=1}^N$, where ξ^i is a noisy measurement of $\mathcal{F}_\xi(\zeta^i)$:

$$\xi^i = \mathcal{F}_\xi(\zeta^i) + \epsilon^i$$

(ϵ^i is the Gaussian noise). The idea is to model the joint distribution of the "input" variable ζ and an "output" variable ξ as a Gaussian Mixture Model. If we join those variables in a vector $\upsilon = [\zeta^T \xi^T]^T$, it is possible to model its probability density function as a mixture of K Gaussian functions

$$p(\upsilon) = \sum_{k=1}^{K} \pi_k \mathcal{N}(\upsilon; \mu_k, \mathbf{\Sigma}_k), \quad \text{such that} \quad \sum_{k=1}^{K} \pi_k = 1$$

where the $\pi_k \in [0\ 1]$ are the priors, and $\mathcal{N}(\upsilon; \mu_k, \mathbf{\Sigma}_k)$ is a Gaussian function with mean μ_k and covariance matrix $\mathbf{\Sigma}_k$:

$$\mathcal{N}(\upsilon; \mu_k, \mathbf{\Sigma}_k) = \left((2\pi)^d |\mathbf{\Sigma}_k|\right)^{-\frac{1}{2}} \exp\left(-\frac{1}{2}(\upsilon - \mu_k)^T \mathbf{\Sigma}_k^{-1}(\upsilon - \mu_k)\right),$$

where d is the dimensionality of the vector υ. The mean vectors μ_k and covariance matrices $\mathbf{\Sigma}_k$ can be separated into their respective input and output components:

$$\mu_k = [\mu_{k,\zeta}^T \ \mu_{k,\xi}^T]^T \qquad \mathbf{\Sigma}_k = \begin{pmatrix} \mathbf{\Sigma}_{k,\zeta} & \mathbf{\Sigma}_{k,\zeta\xi} \\ \mathbf{\Sigma}_{k,\xi\zeta} & \mathbf{\Sigma}_{k,\xi} \end{pmatrix}$$

The Gaussian Mixture Model (GMM) is trained using a standard E-M algorithm, taking the demonstrations as training data. The GMM computes a joint probability density function for the input and the output, so that the probability of the output conditioned on the input are GMM. Hence, it is possible, after training, to recover the expected output variable $\tilde{\xi}$, given the observed input variable ζ.

$$\tilde{\xi} = \tilde{\mathcal{F}}_\xi(\zeta) = \sum_{k=1}^{K} h_k(\zeta)\left(\mu_{k,\xi} + \mathbf{\Sigma}_{k,\xi\zeta}\mathbf{\Sigma}_{k,\zeta}^{-1}(\zeta - \mu_{k,\zeta})\right), \tag{1}$$

where the $h_k(\zeta)$ are given by:

$$h_k(\zeta) = \frac{\pi_k \mathcal{N}(\zeta; \mu_{k,\zeta}, \mathbf{\Sigma}_{k,\zeta})}{\sum_{k=1}^{K} \pi_k \mathcal{N}(\zeta; \mu_{k,\zeta}, \mathbf{\Sigma}_{k,\zeta})}.$$

The tilde ($\tilde{\ }$) sign indicates that we are dealing with expectation values.

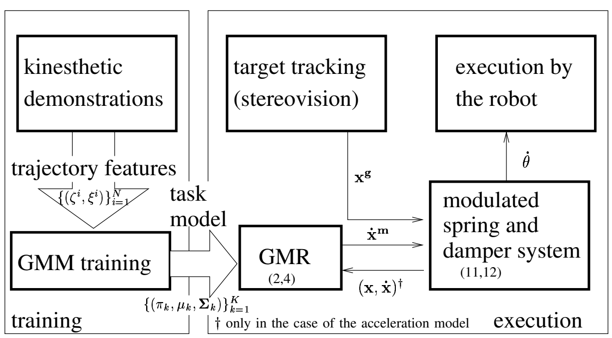

Fig. 1. The architecture of the system. During training the relevant variables (end-effector's position, velocity and acceleration) are extracted from the demonstrations and used to train a GMM. During task execution, this model is used to modulate a spring-and-damper system. $\dot{\mathbf{x}}^\mathbf{m}$ is the end-effector velocity specified by the task model. $\mathbf{x^g}$ is the target location, and $\mathbf{x}^*, \dot{\mathbf{x}}^*, \dot{\theta}^*$ are respectively the actual current end-effector's position and velocity and the joint angles' velocities. The numbers in parentheses refer to the corresponding equations in the text.

C. Acceleration Model

A second way of encoding a trajectory is to take as input the position $\mathbf{x}$ and velocity $\dot{\mathbf{x}}$, and as output the acceleration $\ddot{\mathbf{x}}$. The rationale of this is to consider a trajectory not as a function of time, but as the realization of a second-order dynamical system of the form:

$$\ddot{\mathbf{x}}^\mathbf{m} = \tilde{\mathcal{F}}_{\ddot{\mathbf{x}}}(\mathbf{x}, \dot{\mathbf{x}}). \tag{3}$$

Again, $\tilde{\mathcal{F}}_{\ddot{\mathbf{x}}}$ is obtained by applying (1) with the appropriate variables. The velocity specified by the acceleration model is then given by

$$\dot{\mathbf{x}}^\mathbf{m} = \dot{\mathbf{x}} + \tau\tilde{\mathcal{F}}_{\ddot{\mathbf{x}}}(\mathbf{x}, \dot{\mathbf{x}}), \tag{4}$$

where τ is the time integration constant (set to 1 in this paper). Since the position $\mathbf{x}$ and velocity $\dot{\mathbf{x}}$ depend on the acceleration $\ddot{\mathbf{x}}$ at previous times, this representation introduces a feedback loop, which is not present in the representation given by (2).

III. MODULATED SPRING-AND-DAMPER SYSTEM

We now show how the task model described above is used to modulate a spring-and-damper dynamical system in order to enable a (possibly redundant) robotic arm with n joints to reproduce the task with sufficient flexibility. Although the modulation $\dot{\mathbf{x}}^\mathbf{m}$ is in end-effector space, it is advantageous (for avoiding singularity problems related to inverse kinematics of redundant manipulators) to consider the spring-and-damper dynamical system in joint angle variables:

$$\ddot{\theta}^\mathbf{s} = \alpha(-\dot{\theta} + \beta(\theta^\mathbf{g} - \theta)) \tag{5}$$

where $\theta \in \mathbb{R}^n$ is the vector of joint angles (or arm configuration vector). This dynamical system produces straight paths (in joint space) to the target $\theta^\mathbf{g}$, which acts as an attractor of the system. This guarantees that the robot reaches the target smoothly, despite possible perturbations.

The above dynamical system is modulated by the variable $\dot{\mathbf{x}}^\mathbf{m}$ given by the task model (2) or (4). In order to weigh the modulation, we introduce a modulation factor $\gamma \in \mathbb{R}_{[0\ 1]}$, which weighs the importance of the task model relatively to the spring-and-damper system. If $\gamma = 0$, only the spring-and-damper system is considered, and when $\gamma = 1$ only the task model is considered. In order to guarantee the convergence of the system to $\theta^\mathbf{g}$, γ has to tend to zero at the end of the movement. In the experiments described here, γ is given by:

$$\ddot{\gamma} = \alpha_\gamma(-\dot{\gamma} - \frac{1}{4}\alpha_\gamma\gamma) \quad \text{with } \gamma_0 = 1, \tag{6}$$

where γ_0 is the initial value of γ and $\alpha_\gamma \in \mathbb{R}_{[0\ 1]}$ is a scalar.

Since $\dot{\mathbf{x}}^\mathbf{m}$ lives in the end-effector space (and not in the joint space), the modulation is performed by solving the following constrained optimization problem.

$$\begin{aligned} \dot{\theta} = \quad \underset{\dot{\theta}}{\text{argmin}} \quad & (1-\gamma)(\dot{\theta} - \dot{\theta}^\mathbf{s})^T \bar{\mathbf{W}}_\theta (\dot{\theta} - \dot{\theta}^\mathbf{s}) + \\ & \gamma(\dot{\mathbf{x}} - \dot{\mathbf{x}}^\mathbf{m})^T \bar{\mathbf{W}}_\mathbf{x} (\dot{\mathbf{x}} - \dot{\mathbf{x}}^\mathbf{m}) && (7) \\ \text{u.c.} \quad & \dot{\mathbf{x}} = \mathbf{J}\dot{\theta}, && (8) \end{aligned}$$

where $\mathbf{J}$ is the Jacobian of the robot arm kinematic function $\mathbf{K}$ and $\bar{\mathbf{W}}_\theta \in \mathbb{R}^{n\times n}$ and $\bar{\mathbf{W}}_\mathbf{x} \in \mathbb{R}^{m\times m}$ are diagonal matrices necessary to compensate for the different scale of the $\mathbf{x}$ and θ variables. As a rough approximation, the diagonal elements of $\bar{\mathbf{W}}_\mathbf{x}$ are set to one and those of $\bar{\mathbf{W}}_\theta$ are set to the average

distance between the robot base and its end-effector.
The solution to this minimization problem is given by [12]:

$$\dot{\theta} = \left(\mathbf{W}_\theta + \mathbf{J}^T\mathbf{W}_\mathbf{x}\mathbf{J}\right)^{-1}\left(\mathbf{W}_\theta\dot{\theta}^\mathbf{s} + \mathbf{J}^T\mathbf{W}_\mathbf{x}\dot{\mathbf{x}}^\mathbf{m}\right) \quad (9)$$

$$\text{where} \quad \mathbf{W}_\theta = (1-\gamma)\bar{\mathbf{W}}_\theta, \quad \mathbf{W}_\mathbf{x} = \gamma\bar{\mathbf{W}}_\mathbf{x}. \quad (10)$$

To summarize, the task is performed by integrating the following dynamical system:

$$\ddot{\theta}^\mathbf{s} = \alpha(-\dot{\theta} + \beta(\theta^\mathbf{g} - \theta)) \quad (11)$$

$$\dot{\theta} = \left(\mathbf{W}_\theta + \mathbf{J}^T\mathbf{W}_\mathbf{x}\mathbf{J}\right)^{-1}\left(\mathbf{W}_\theta\dot{\theta}^\mathbf{s} + \mathbf{J}^T\mathbf{W}_\mathbf{x}\dot{\mathbf{x}}^\mathbf{m}\right) \quad (12)$$

where $\mathbf{W}_\mathbf{x}$ and $\mathbf{W}_\theta$ are given by (6) and (10), and $\dot{\mathbf{x}}^\mathbf{m}$ is given either by (2) (velocity model) or by (4) (acceleration model). Integration is performed using a first-order Newton approximation ($\dot{\theta}^\mathbf{s} = \dot{\theta} + \tau\ddot{\theta}^\mathbf{s}$).

Since the target location is given in cartesian coordinates, inverse kinematics must be performed in order to obtain the corresponding target joint angle configuration which will serve as input of the spring-and-damper dynamical system. In the case of a redundant manipulator (such as the robot arm used in the following experiments) the desired redundant parameters of the target joint angle configuration can be extracted from the demonstrations. This is done by using the GMR technique described in Table I to build a model of the final arm configuration as a function of the target location.

Using an attractor system in joint angle space has the practical advantage of reducing the usual problems related to end-effector control, such as joint limit and singularity avoidance. Equation 9, which is a generalized version of the Damped Least Squares inverse [13] [14], is a way to simultaneously control the joint angles and the end-effector, imposing soft constraints on both of them. It is thus different than optimizing the joint angles in the null space of the kinematic function.

IV. Experiments

A. Setup

We validate and compare the systems described in this paper on two experiments. The first experiment involves a robot putting an object into a box and the second experiment consists in reaching and grasping for an object. Those experiments were chosen because (1) they can be considered as simple goal-directed tasks (for which the system is intended), (2) they are tasks commonly performed in human environments and (3) they presents a clear success or failure criterion.
All the experiments presented below are performed with a Hoap3 humanoid robot acquired from Fujitsu. This robot has four back-drivable degrees of freedom (dof) at each arm. Thus, the robot arms are redundant, as we do not consider end-effector orientation. The robot is endowed with a stereo-vision system enabling it to track color blobs. A small color patch is fixed on the box and on the object to be grasped, enabling their 3D localization. Pictures of the setup are shown in Fig. 2.

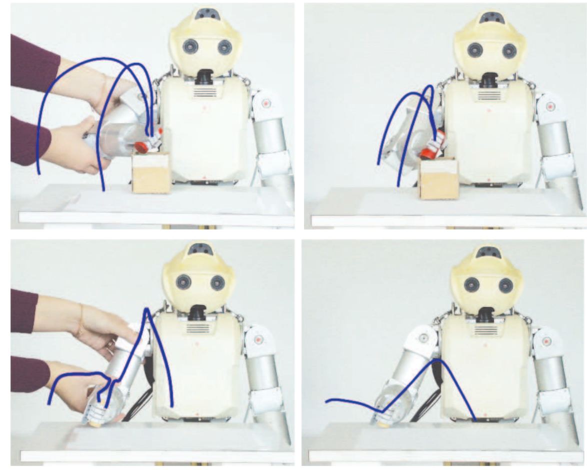

Fig. 2. The setup of the experiments. The top pictures show the first task and the lower picture sow the second task Left: a human operator demonstrates a task to the robot by guiding its limbs. Right: the robot performs the task, starting from different initial positions.

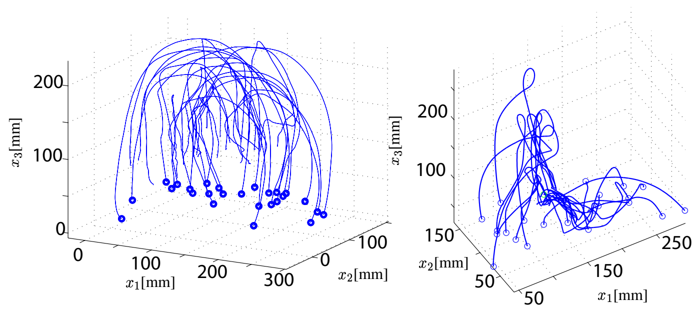

Fig. 3. The demonstrated trajectories for the box task (left) and the grasping task (right). Circles indicate starting positions.

1) Preprocessing: During the demonstrations, the robot joint angles were recorded and the end-effector positions were computed using the arm kinematic function. All recorded trajectories were linearly normalized in time ($T = 500$ time steps) and Gaussian-filtered to remove noise. The number of Gaussian components for the task models were found using the Bayesian Information Criteria (BIC) [15], and the parameter values used were $\alpha_\gamma = 0.06$, $\alpha = 0.12$ and $\beta = 0.06$.

B. Putting an object into a box

1) Description: For this task, the robot is taught to put an object into the box (see Fig.2). In order to accomplish the task, the robot has to avoid hitting the box while performing the movement and must thus first reach up above the box and then down to the box. A straight line reaching will in general cause the robot to hit the box while reaching and thus fail.

2) Training: A set of 26 kinesthetic demonstrations were performed, with different initial positions and box locations. The box was placed on a little table. Thus its location only varies in the horizontal plane. Similarly, the initial position of the object (and thus of the end-effector) lied on the table. The set of demonstrated trajectories is depicted in Fig. 3, left. The velocity models trained on this data are shown in Fig. 4, left.

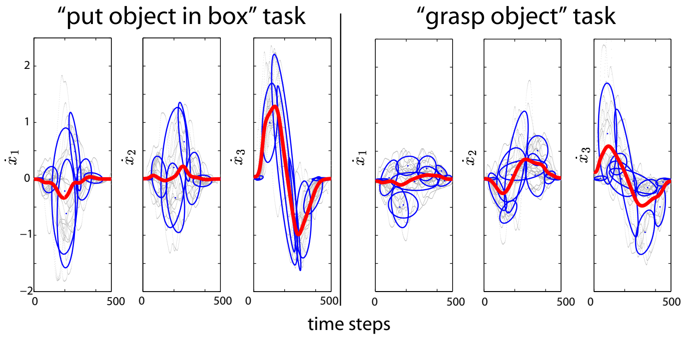

Fig. 4. The velocity models for both tasks. The dots represent the training data, the ellipses the Gaussian components and the thick lines the trajectory obtained by GMR alone. The thick lines show that, for the first task, the horizontal components $\dot{x}_1$ and $\dot{x}_2$ are averaged out by the model, but the vertical component $\dot{x}_3$ shows a marked upward movement. For the second task, all components are almost averaged out.

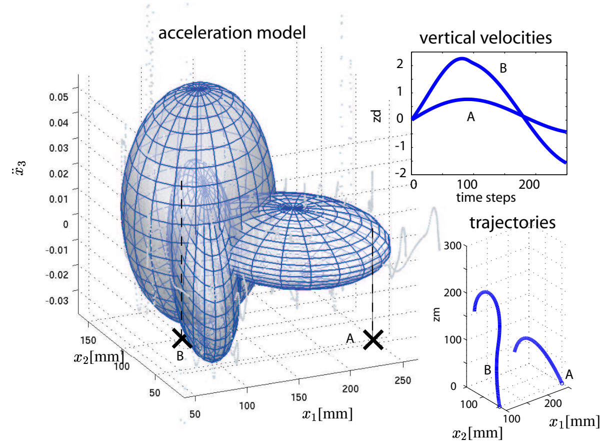

Fig. 5. In the center, the acceleration model for the second task. The ellipsoids show the Gaussian components at twice their standard deviation. Only three projections (out of nine) are shown. The vertical acceleration strongly depends on the position in the horizontal plane. On the lower right, two trajectories encoded by this model but starting from different positions A and B (indicated by the crosses) are shown. The corresponding vertical velocity profiles appear on the upper right. They differ significantly, as the model is not homogeneous across the horizontal plane.

C. Reach and Grasp

1) Description: In order to accomplish this task, the robot has to reach and correctly place its hand to grasp a chess piece. In other words it has to place its hand so that the chess piece stands between its thumb and its remaining fingers, as shown in Fig. 7, left. This figure illustrates that the approaching the object can only be done in one of two directions: downward or forward. This task is more difficult than the previous one, as the movement is more constrained. Moreover, a higher precision is required on the final position, since the hand is relatively small.

2) Training: A set of 24 demonstrations were performed starting from different initial positions located on the horizontal plane of the table. The chess piece remained in a fixed location. Depending on the initial position, the chess piece was approached either downward or forward (as illustrated on Fig. 7). The set of demonstrations is represented in Fig. 3, right. The resulting velocity model is shown in Fig. 4, right. One can notice that there is no velocity feature that is common to all demonstrated trajectories. The acceleration model is shown in Fig. 5. This model captures well the fact that the vertical acceleration component depends on the position in the horizontal plane.

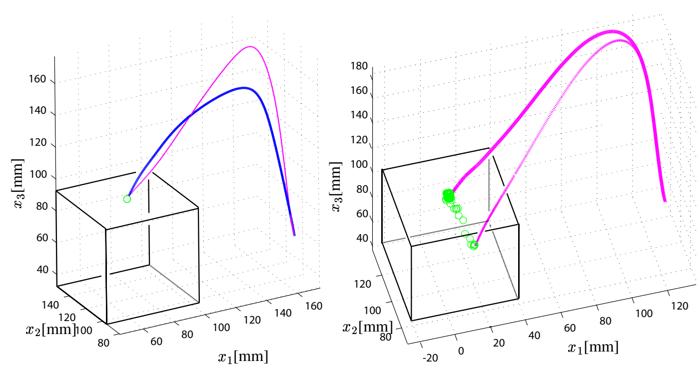

Fig. 6. Left: end-effector trajectories of the robot putting the object into the box. The thin line corresponds to the velocity model and the thick line corresponds to the acceleration model. Right: online trajectory adaptation to a target displacement using the velocity model. The circles indicate to location of the box, as tracked by the stereo-vision system. The thick line shows the produced trajectory and the thin line shows the original trajectory if the box remained unmoved. Similar results were obtained with the acceleration model.

D. Results

Endowed with the system described above, the robot is able to successfully perform both tasks. For the first task, both the velocity and the acceleration models can produce adequate trajectories (see Fig. 6, left for examples). The system can adapt its trajectory online if the box is moved during movement execution (see Fig. 6, right). For the second task, examples of resulting trajectories are displayed in Fig. 7, right. In order to evaluate the generalization abilities of the systems, both tasks were executed from various different initial positions arbitrarily chosen on the horizontal plane of the table, and covering the space reachable by the robot. Fig. 8 shows the results and starting positions for both experiments. For the box experiment (left), the velocity model was successful for 22 out of the 24 starting locations (91%). The two unsuccessful trials, indicated by empty circles, correspond to initial positions close to the work space boundaries. The acceleration model was successful for all trials (100%).

For the chess piece experiment (Fig. 8, right), the velocity model was successful for 5 out of 21 (24%) trials whereas the acceleration model was successful for 18 trials (86%). This performance gap is due to the fact that this task does not require a fixed velocity modulation. The adequate modulation depends on the position. This position-dependent modulation can be captured by the acceleration model, but not by the velocity model. As illustrated in Fig. 5, the acceleration model is able to produce different velocity profiles, depending on the starting position and is thus more versatile than the velocity model.

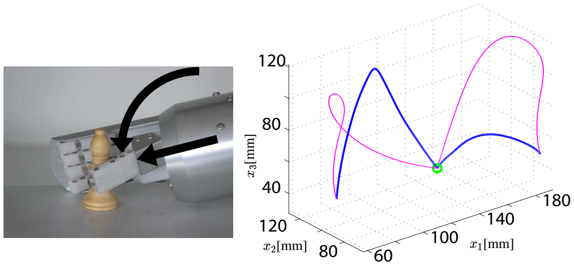

Fig. 7. Left: the chess piece to be grasped. For a successful grasp, the robot has to approach it as indicated by the arrows. Right: resulting trajectories for the grasping task, starting from two different initial positions. The acceleration model (thick lines) adapts the modulation to the initial position, while the velocity model (thin lines) starts upward in both cases. The trajectory produced by the velocity model and starting left of the target is not successful.

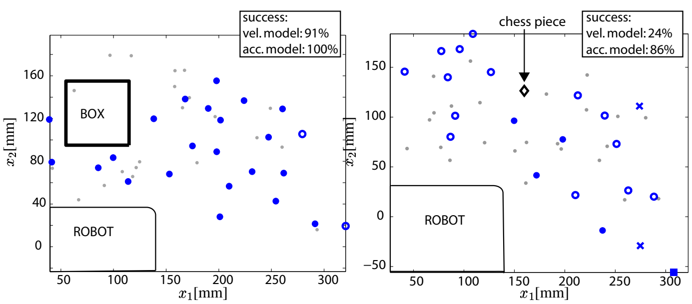

Fig. 8. The robustness to initial end-effector position for both tasks. The plots represent top views of the first (right) and second (left) experiment. The filled markers (circles or squares) indicate all initial positions for which the velocity model was successful. The circles (filled and non-filled) indicate all initial positions for which the acceleration model was successful. The crosses indicate initial end-effector position, for which both models failed. The dots indicate the starting positions of the training set.

V. Discussion

Our results show that the framework suggested in this paper can enable a robot to learn constrained reaching tasks from kinesthetic demonstrations, and generalize them to different initial conditions. Using a dynamical system approach allows to deal with perturbations occurring during the task execution. This framework can be used with various task models and has been tested for two of them, the velocity model and the acceleration model. The results indicate that the velocity model is too simplistic if the task requires different velocity profiles when starting from different positions in the workspace. The acceleration model is more sophisticated and can model more constrained movements, but may fail to provide an adequate trajectory when brought away from the demonstrations in the phase space $(\mathbf{x}, \dot{\mathbf{x}})$. Other regressions techniques, such as LWR, could also be used. But if there are inconstancies across demonstrations, simple averaging may fail to provide adequate solutions.

In its present form, the modulation factor between the dynamical system and the task model (γ) is not learned. Learning it from the demonstrations is likely to further improve the performance of the system, especially for tasks requiring a modulation at the end of the movement. It would also be desirable to have a system that extracts the relevant variables, and automatically selects the adequate model. A first step in this direction has been taken in [16], where a balance between different sets of variables is achieved.

Of course, the adequacy of this framework is restricted to relatively simple tasks, such as those described in the experiments. More complicated tasks, such as obstacle avoidance in complex environments or stable grasping of particular objects require a detailed model of the environment and more elaborate planning techniques. The tasks considered for this framework are those that cannot be accomplished by simple point-to-point reaching, but still simple enough to avoid the complete knowledge of the environment. But this framework could be extended to learn more complicated tasks. In a first step in this direction, [17] investigates in simulation and on a simplified framework how Reinforcement Learning can deal with obstacle avoidance.

References

[1] A. Billard and R. Siegwart, Eds., *Robotics and Autonomous Systems, Special Issue: Robot Learning From Demonstration*. Elsevier, 2004, vol. 47, no. 2,3.

[2] A. Ijspeert, J. Nakanishi, and S. Schaal, "Movement imitation with nonlinear dynamical systems in humanoid robots," in *Proceedings of the IEEE International Conference on Robotics and Automation*, 2002, pp. 1398–1403.

[3] W. Ilg, G. Bakir, J. Mezger, and M. Giese, "On the representation, learning and transfer of spatio-temporal movement characteristics," *International Journal of Humanoid Robotics*, pp. 613–636, 2004.

[4] A. Shon, J. Storz, and R. Rao, "Towards a real-time bayesian imitation system for a humanoid robot," in *Proceedings of the IEEE International Conference on Robotics and Automation*, 2007, pp. 2847–2852.

[5] O. Jenkins, G. González, and M. Loper, "Tracking human motion and actions for interactive robots," in *Proceedings of the Conference on Human-Robot Interaction (HRI07)*, 2007, pp. 365–372.

[6] D. Grimes, D. Rashid, and R. Rao, "Learning nonparametric models for probabilistic imitation," in *Advances in Neural Information Processing Systems (NIPS 06)*, 2006.

[7] K. Ogawara, J. Takamatsu, H. Kimura, and K. Ikeuchi, "Extraction of essential interactions through multiple observations of human demonstrations," *IEEE Trans. Ind. Electron.*, vol. 50, pp. 667–675, 2003.

[8] R. Dillmann, "Teaching and learning of robot tasks via observation of human performance," *Robotics and Autonomous Systems*, no. 2,3, pp. 109–116, 2004.

[9] C. Campbell, R. Peters, R. Bodenheimer, W. Bluethmann, E. Huber, and R. Ambrose, "Superpositioning of behaviors learned through teleoperation," *IEEE Transactions on Robotics*, 2006.

[10] R. Burridge, A. Rizzi, and D. Koditschek, "Sequential composition of dynamically dexterous robot behaviors," *International Journal of Robotics Research*, 1999.

[11] Z. Ghahramani and M. Jordan, "Supervised learning from incomplete data via an em approach," in *Advances in Neural Information Processing Systems 6*, J. Cowan, G. Tesauro, and J.Alspector, Eds. Morgan Kaufmann Publishers, 1994.

[12] A. Billard, S. Calinon, and F. Guenter, "Discriminative and adaptive imitation in uni-manual and bi-manual tasks," *Robotics and Autonomous Systems*, vol. 54, no. 5, pp. 370–384, 2006.

[13] C. Wampler, "Manipulator inverse kinematic solutions based on vector formulations and damped least-squares methods," *IEEE Transactions on Systems, Man and Cybernetics, Part C*, vol. 16, no. 1, pp. 93–101, 1986.

[14] Y. Nakamura and H. Hanafusa, "Inverse kinematics solutions with singularity robustness for robot manipulator control," *ASME Journal of Dynamic Systems, Measurement, and Control*, vol. 108, pp. 163–171, 1986.

[15] G. Schwarz, "Estimating the dimension of a model," *Annals of Statistics*, vol. 6, 1978.

[16] S. Calinon, F. Guenter, and A. Billard, "On learning, representing and generalizing a task in a humanoid robot," *IEEE Trans. Syst., Man, Cybern. B*, vol. 37, no. 2, pp. 286–298, 2007.

[17] F. Guenter, M. Hersch, S. Calinon, and A. Billard, "Reinforcement learning for imitating constrained reaching movements," *RSJ Advanced Robotics*, vol. 21, no. 13, pp. 1521–1544, 2007.

Abstract Archive: 2000-2005

printer-friendly version

Home » Publications » Abstracts/Annual Meeting Publications » Abstract Archive: 2000-2005

Neuroscience 2005 Abstract

Presentation Number:	469.7
Abstract Title:	Eye to eye communication.
Authors:	**Rustichini, A.***[1]; Fadiga, L.[2]; Lungu, O.[3] [1]Economics, Univ. of Minnesota, Minneapolis, MN [2]Italy, 1035 Heller Hall, 71, 19th Avnue South inneapolis, MN 55455 , 55455, [3]SBTA, Sezione di Fisiologia Umana, 1035 Heller Hall, 71, 19th Avnue South inneapolis, MN 55455 , 55455,
Primary Theme and Topics	Cognition and Behavior - Human Cognition, Behavior, and Anatomy -- Perception and imagery
Session:	469. Perception and Imagery I Slide
Presentation Time:	Monday November 14, 2005 2:30 PM-2:45 PM
Location:	Washington Convention Center - Room 143A
Keywords:	Sympathy, Mirror Systems

Sympathy is the ability of the observer to reproduce the internal states of others, either when observing an external event or the display of a reaction, motor or affective. We test the hypothesis that sympathy is used as an information extracting device: the reproduction of the neural activity of the observed subject provides a signal on the information available to the observed subject. An implication of the theory is that a subject has very little to know on his own internal states, so brain activity related to sympathy should be smaller than it is when a different subject is involved. We test this hypothesis using the simplest form of interpersonal communication: the exchange of gazes among human subjects, including the subject looking at himself. Five different conditions have been used. The key comparisons are between the brain activity of a subject when he is looking at a different person and when he is looking at his own eyes. In other conditions, subjects are looking at an observer who is not looking, or they are looked at as they are not looking.
A group of 29 subjects has been observed in an fMRI study. The results support the hypothesis of sympathy ax an information acquisition. For example, BA 44 is involved specifically when two subjects exchange gazes. Anterior Insula is activated when subjects are being looked at and are not looking.

Supported by NSF

Sample Citation:

[Authors]. [Abstract Title]. Program No. XXX.XX. 2005 Neuroscience Meeting Planner. Washington, DC: Society for Neuroscience, 2005. Online.

Return to the Previous Page
Home

Brain Advance Access published May 14, 2009

BRAIN
A JOURNAL OF NEUROLOGY

Encoding of human action in Broca's area

Patrik Fazio,[1,2] Anna Cantagallo,[2] Laila Craighero,[1] Alessandro D'Ausilio,[1] Alice C. Roy,[3] Thierry Pozzo,[4,5] Ferdinando Calzolari,[2] Enrico Granieri[2] and Luciano Fadiga[1,5]

1 DSBTA, Section of Human Physiology, University of Ferrara, Ferrara, Italy
2 Department of Neuroscience Rehabilitation, Hospital and University of Ferrara, Ferrara, Italy
3 Institute of Cognitive Science, CNRS, Lyon, France
4 INSERM-ERM 207 Motricité Plasticité, Université de Bourgogne, Dijon, France
5 Italian Institute of Technology (IIT), Genova, Italy

Correspondence to: Prof. Luciano Fadiga, DSBTA,
Section of Human Physiology,
University of Ferrara,
Via Fossato di Mortara 17/19,
44100 Ferrara, Italy
E-mail: luciano.fadiga@unife.it

Broca's area has been considered, for over a century, as the brain centre responsible for speech production. Modern neuroimaging and neuropsychological evidence have suggested a wider functional role is played by this area. In addition to the evidence that it is involved in syntactical analysis, mathematical calculation and music processing, it has recently been shown that Broca's area may play some role in language comprehension and, more generally, in understanding actions of other individuals. As shown by functional magnetic resonance imaging, Broca's area is one of the cortical areas activated by hand/mouth action observation and it has been proposed that it may form a crucial node of a human mirror-neuron system. If, on the one hand, neuroimaging studies use a correlational approach which cannot offer a final proof for such claims, available neuropsychological data fail to offer a conclusive demonstration for two main reasons: (i) they use tasks taxing both language and action systems; and (ii) they rarely consider the possibility that Broca's aphasics may also be affected by some form of apraxia. We administered a novel action comprehension test—with almost no linguistic requirements—on selected frontal aphasic patients lacking apraxic symptoms. Patients, as well as matched controls, were shown short movies of human actions or of physical events. Their task consisted of ordering, in a temporal sequence, four pictures taken from each movie and randomly presented on the computer screen. Patient's performance showed a specific dissociation in their ability to re-order pictures of human actions (impaired) with respect to physical events (spared). Our study provides a demonstration that frontal aphasics, not affected by apraxia, are specifically impaired in their capability to correctly encode observed human actions.

Keywords: Broca's area; action recognition; mirror-neuron system; frontal aphasia; motor syntax

Abbreviations: IFG = inferior frontal gyrus; LST = Language sequencing task; MRI = magnetic resonance imaging; RT = reaction time; TT = trial time

Introduction

The seminal work of the French neurologist Paul Broca established that the posterior part of the left inferior frontal gyrus (IFG) was of critical importance for speech production. Broca's famous case, Leborgne, suffered from left frontal damage extending from the inferior part of the third frontal circumvolution to parts of the insula and the striatum (Broca, 1861; Dronkers *et al*., 2007). Broca's aphasia was thus described as a syndrome characterized by effortful speech production, impairment in melodic line and

Received November 12, 2008. Revised January 29, 2009. Accepted March 30, 2009

articulation, semantic and phonemic paraphasias, telegraphic sentences with reduced and abnormal grammatical forms (Broca, 1861; Alexander *et al*., 1990; Caplan *et al*., 1996).

The first empirical evidence that Broca's area is involved in speech production was provided by Penfield and Roberts (1959). These authors demonstrated that the electrical stimulation of Broca's area in awake neurosurgery patients could evoke a complete arrest of ongoing speech. The hot spot for this effect was located in the *pars opercularis* of the IFG (see also Ojemann *et al*., 1989). Moreover, Dronkers (1996) showed that a lesion affecting the most posterior part of left IFG (involving insula as well) lead to apraxia of speech (AOS). AOS deficit can be defined as a disorder in the motor programming of the speech musculature to produce the correct sound of words in the proper sequence with the appropriate timing.

Recently, a more complex picture of the role played by Broca's area in the language domain has been given. Several studies demonstrated that Broca's aphasics, in addition to their deficits in production, are also impaired in speech comprehension. Deficits are more evident when patients were tested with verbal material requiring syntactical understanding (Caramazza and Zurif, 1974; Alexander *et al*., 1990; Caplan *et al*., 1996). The role of Broca's area in understanding speech has been further supported by the work by Schäffler and collaborators (1993, 1996) showing that the electrical stimulation of Broca's area in non-aphasic neurosurgery patients may elicit comprehension deficits of complex verbal commands.

Language-related studies aside, several recent works have found activation of Broca's area in other cognitive domains (for a review see Fadiga *et al*., 2006) and, more interestingly as far as the objectives of the present study are concerned, in action viewing, action execution and action imitation (Grafton *et al*., 1996; Binkofski *et al*., 1999; Iacoboni *et al*., 1999; Nishitani and Hari, 2000; Buccino *et al*., 2001; Grèzes and Decety 2001; Baumgaertner *et al*., 2007). These data have been considered as an empirical support to the existence of a mirror-like system in humans, mapping execution and observation of actions onto the same neural substrate (Rizzolatti and Craighero, 2004). However, some concerns have been raised regarding the conclusions that can be drawn from such techniques and experimental designs (Dinstein *et al*., 2007; Turella *et al*., 2008). In fact, the use of a correlational approach cannot provide a final proof of the involvement of Broca's area in the human mirror-neuron system.

A possible answer to the question whether Broca's area could be involved in action understanding might be provided both by neuropsychological studies of brain lesioned patients and by temporary inactivation of Broca's area by transcranial magnetic stimulation (TMS) during action-understanding tasks. Pobric and colleagues (2006) administered TMS on the IFG while subjects had to judge the weight of an object lifted by an actor. Their data show a reduced accuracy in performing the task, in accordance with the hypothesis that the IFG plays an important role in encoding the details of action kinematics. Moreover, several studies on frontal aphasic patients have shown a correlation between lesion location and action-related non-verbal impairments such as recognizing signs, gestures and pantomimes (Duffy and Duffy, 1975; Gainotti and Lemmo, 1976; Daniloff *et al*., 1982; Varney, 1982; Glosser *et al*., 1986; Wang and Goodglass, 1992; Bell, 1994). More recently, Tranel *et al*. (2003), showed that left frontal brain-damaged patients have difficulty in understanding action details when presented with cards depicting various actions, and Saygin *et al*. (2004) have demonstrated a significant correlation between linguistic deficits and the comprehension of actions in patients with different types of aphasia.

It should be noted, however, that although strongly suggestive of a strict relationship between language- and action-related domains, the results by Tranel (2003) and Saygin (2004) were achieved through tasks including some linguistic components (i.e. verbal instructions). Therefore, the reported deficits in the action domain might have, at least partially, been altered by uncontrolled linguistic processes. Furthermore, left fronto-parietal lesions are often associated with praxic disturbances (Goldenberg, 1996). Both studies, unfortunately, did not control for such possibilities, which could act as a critical confounding factor in the light of the recent study by Pazzaglia *et al*. (2008), showing that limb apraxic deficits are often associated to the impairment of gesture comprehension.

As a consequence, on the basis of the current empirical knowledge, a conclusive picture of the causal relationship between Broca's aphasia and action understanding deficits cannot be drawn without doubts. In the present work, we selected frontal aphasic patients (without apraxia) on the basis of lesion localization. We then administered a newly designed task to measure patient's performance in action comprehension without taxing the language system. Patients were requested to correctly sequence some randomly mixed pictures taken from video clips representing human actions or physical events. Our prediction was that Broca's aphasics would exhibit a dissociation in dealing with these two classes of stimuli, thus providing the evidence that Broca's area, beside its linguistic function, is also involved in encoding human actions.

Methods

Participants

Medical records of twenty patients from the community of Ferrara (Department of Neuroscience, University and Hospital of Ferrara, Unit of Neuropsychological Rehabilitation, Italy) were evaluated after obtaining informed consent. Patients were selected if, at the time of the enrolment into the hospital rehabilitation program, they presented a vascular lesion in the territory of the left middle cerebral artery, according to computerized tomography (CT) or magnetic resonance imaging (MRI) data. All of them presented disorders of language production with agrammatic speech (speech was laboured, choppy and poorly articulated), but their comprehension of normal conversation was well preserved. In addition to the testing for aphasia (Ciurli *et al*., 1996; Capasso and Miceli, 2001; Token test: De Renzi and Vignolo, 1962), patients were screened for the presence of apraxia (De Renzi *et al*., 1966; De Renzi *et al*., 1980). Further exclusion criteria included diagnosis or suspicion of dementia, head traumas, brain tumours, multiple infarcts or other neurological conditions. All patients had normal intelligence and had no difficulty in attending to, perceiving or retrieving visual stimuli. According to the evaluation

Table 1 Socio-demographic data and lesion location

Initials	Gender	Age	Education	Main lesions
DF	M	53	8	Frontal (+IFG), temporal, parietal, insula + external capsule region.
FG	M	51	17	Frontal (+IFG), temporal, insula + external capsule region and basal ganglia.
SC	F	27	18	Frontal (+IFG), temporal, insula + external capsule region.
EC	M	56	8	Frontal (+IFG), temporal, insula.
CC	M	61	13	Frontal (+IFG), temporal, insula + external capsule region.
GF	M	60	8	Frontal (+IFG), parietal, insula.

Age, gender and lesion locations of patients that fulfilled the experimental requirements.

of their case history, 11 patients out of the initial 20 were considered for further testing.

The further neuropsychological testing (see below) was aimed at selecting patients with a high degree of cognitive functionality and with normal praxic capabilities. Following this second-level, more restrictive testing, the number of patients recruited for the final experimental phase was reduced from eleven to six. The age of recruited patients ranged from 27 to 61 (mean 51 ± 12.5 SD) and their average level of education was 12 ± 4.7 years. Socio-demographic data and lesion location, as assessed by the local neuroradiology unit, are provided in Table 1. At the time of our investigation, these patients presented a stable lesion due to cerebro-vascular accident, which had occurred 3–6 years before the enrolment, and none of them had been included in other studies. As a control group we selected six adult participants matched for age, handedness and education level, with no history of neurological or psychiatric disorders. All of them had normal or corrected-to-normal vision and hearing. The control group had a mean age of 50.2 ± 13.1 and a mean educational level of 12.7 ± 4 years. The procedures used in the study were in agreement with the guidelines of the University of Ferrara Ethical Committee and with the Declaration of Helsinki.

Neuropsychological testing

Patients were tested by a skilled neurologist in a quiet room reserved for experimental purposes, in the neuropsychological rehabilitation unit. We collected a series of reduced versions of standard neuropsychological questionnaires aimed at specifically and rapidly testing a wider range of cognitive functions. This test was administered both to the eleven patients and to the control subjects. The main goal of this procedure was to select patients with a high degree of cognitive functionality. The test included three main sections—one testing general abilities, the second concerning praxis and the third related to the language domain. The first section included items evaluating calculus, memory span and rhythm generation as well as the general sense of direction (questions such as: 'Where are we now?', 'Why we are here?', 'What day of the week is today?'). The praxic section had 29 items with a cut-off level of 18 correct responses. The testing included: (i) imitation of distal intransitive movements; (ii) imitation of intransitive movements of the mouth area; (iii) imitation of transitive movements; (iv) execution of intransitive sequence of movements upon verbal instruction; and (v) pantomime. The language section included 67 items with a cut-off level of 18 correct responses. The complete set included: (i) denomination of visually presented natural and manufactured objects and tools; (ii) repetition of words and pseudo-words following audio-visual and auditory presentation; (iii) verbal fluency; and (iv) auditory comprehension testing. This collection of tests was administered prior to the experimental session and only six patients met the second enrolment criteria for participating in the study. The whole experimental session was videotaped for further offline analysis. We then visually inspected each movie in order to exclude the presence of any sub-clinical apraxic signs. Particular attention was devoted to the exclusion of deficits regarding the temporal and spatial sequencing of an action or the loss of object knowledge. This offline analysis was carried out by two independent professional neuropsychologists. The outcome of this additional evaluation confirmed previous examinations.

Lesion analysis

To anatomically characterize the brain lesion, all recruited patients underwent an additional specific MRI session at the beginning of the study. MRI images were acquired through three-dimensional-fast spoiled gradient recalled (3D-FSPGR) T_1-weighted sequence (TR 12.6 ms, TE 2.7 ms, TI 400 ms, FOV 250 mm $\times$ 250 mm, thickness 0.6 mm, gap: 0.6 mm, 256×256 matrix, 250 slices). The graphical outline of the lesions and the co-registration of individual brains in standard stereotaxic space were performed offline using MRIcro software (Rorden and Brett, 2000). We transformed each anatomical image into a standard stereotaxic space using the co-registration method (non-linear warping) provided by SPM 96. Brain lesions were mapped in the MNI stereotaxic space using the standard MRI volume redefined by Colin's Atlas (Evans *et al*., 1993). After individual co-registration, the lesioned areas of each patient were superimposed onto each other by means of the specific tool provided by MRIcro software. We thus obtained the region of superimposition common to all the patients shown in red in Fig. 1A. The higher lesion overlap of this region was centred in the *pars opercularis* of the IFG (BA44), as shown by the probabilistic atlas of BA44 by Amunts *et al*. (1999) (Fig. 1B and C).

Experimental design, materials and procedure

Participants sat comfortably in front of a touch screen monitor (MicroTouch M170, 3M Touch Systems, Inc.). They were informed of the experimental procedure and completed a first practice trial under the experimenter's guidance. The experimental interface was run on a PC using custom made software. All experimental events and their relative timing were automatically recorded during each trial and stored on the hard-disk for further offline analysis. Each trial began with a message displayed on the computer screen inviting the participant to press a key to start the trial. A videoclip was then displayed on the screen and the participant was instructed to pay attention to it. At the end of the videoclip, after a delay of 0.5 s with black screen, four images taken from the same video were presented simultaneously at four different spatial locations

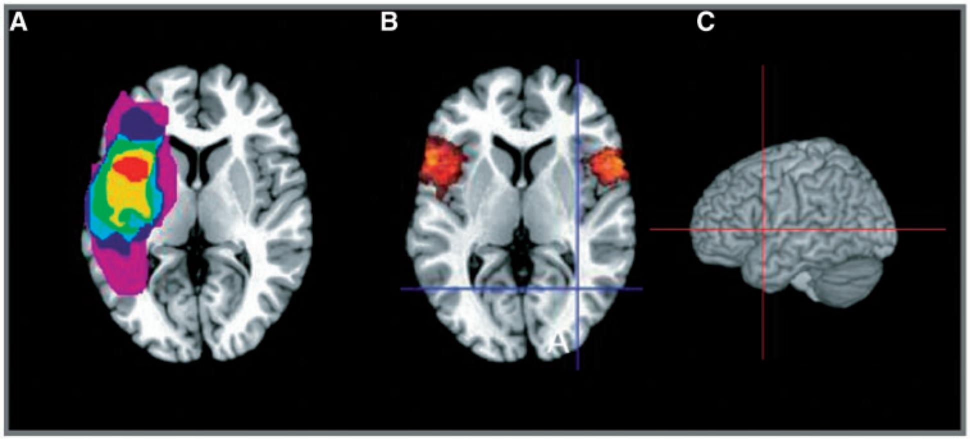

Figure 1 Anatomical location of patient's lesions. (**A**) shows the overlays of the six patients lesions on the Colin's template brain, where different colours represent each patient, (**B**) shows the probabilistic extension of Broca's area as identified by Amunts *et al*. (1999), on the same horizontal slice and (**C**) shows the 3D rendering of the Colin's standard brain with the overlap of lesions marked with a red cross on the surface. Note that the overlap corresponds perfectly with the *pars opercularis* of the IFG (BA44).

(upper left, upper right, lower left and lower right). At this point subjects were asked to organize the four snapshots in a meaningful temporal order, by touching the screen. As soon as two snapshots were touched sequentially, their spatial location on the screen was swapped. When participants considered they had accomplished the trial, they had to press a validation button. All subjects were instructed to be as accurate as possible and, only in the second instance, to complete the task as fast as possible. The particular stress given to accuracy was justified by the observation that some patients had more difficulty when temporally pressed, also depending on the severity of their clinical picture.

Participants were requested to constantly focus their attention on the task and, after each trial, were asked whether they needed a rest. Generic motivational feedback (e.g. 'you are doing great so far', 'very good') was given as often as considered necessary to keep participants engaged in the task (approximately once every trial). Feedback regarding the accuracy of the performance ('OK!' or 'Fail!') was given on the computer screen at the end of each trial. At the end of each trial, subjects were also asked to explain what the video clip was about, to verify that they could understand the global meaning of the stimuli.

The following variables were recorded: accuracy (degree of correct sequencing), reaction time (RT) and trial time (TT). Accuracy measured whether the order of the snapshots, generated by the subjects, was correct or not (in percentage). RT was the amount of time elapsed from presentation of the snapshots to the first touch of the screen. TT was the amount of time between presentation of the snapshots and pressing of the validation button.

Sequencing task

The same task was administered to both patients and controls. The video clips were subdivided into two different classes: human actions and physical events. Human actions stimuli were transitive and intransitive actions performed by a human agent (e.g. hand-grasping of a bottle, head-turning and pointing, etc.). Video clips of physical events represented common life dynamic events, such as a bicycle falling on the floor or a door opening by itself. Snapshots were selected for each video clip by the experimenters, paying attention to provide enough cues for the successive sequencing task (see Fig. 2 for a pictorial description of the task).

Nineteen videos (Table 2), plus one used to familiarize the participants to the experimental procedure, were used during the experiment. In order not to overload the patient's attention, we restricted the number of videos by unbalancing the number of stimuli of the two categories (human actions: 14 and physical events: 5). This decision was taken after a pilot experiment on 13 healthy subjects using a larger set of 30 movies. We found that physical events, on average, led to a smaller variability of the time necessary to accomplish the task (i) and of the time to begin sequencing the four pictures (ii) [standard deviation (SD) (i) 4.52 s; (ii) 4.02 s for human actions and (i) 2.29 s; (ii) 1.68 s for physical events]. However, task difficulty could be better described by absolute time values, rather than SD. In fact, SD indicates how variable performance is across different subjects, whereas mean values describe how difficult the task was in all subjects. Therefore, we reduced the number of items in the least variable condition (physical events) but selected those trials that, according to mean values, were homogeneously spread across the difficulty continuum (Table 2).

Language sequencing task

Patients also underwent a second testing phase. Their task was similar to that outlined previously, but the stimuli differed. The videoclips were replaced with either written sentences (8) or single words (20). Their task was to sequence four scrambled written segments taken from the stimuli. Sentences were divided into simpler constituents (i.e. Press/the button/to open/the door) and words into syllables (i.e. Cam/mi/na/re: to walk).

Statistical analysis

A Mann–Whitney U-test was used to analyse the neuropsychological data obtained from both patients and normal controls for each sub test. The null hypothesis was that the two samples are drawn from a single population, and therefore their performance is similar. Three two-way repeated-measure ANOVAs were performed on RT, TT and Accuracy, with a between-subjects factor GROUP (Aphasics, Controls) and a within-subject factor CONDITION (human actions, physical events). Fisher's LSD *post hoc* comparisons were then conducted when factors showed a significant effect. Furthermore, a linear correlation analysis was run between patient performance in the

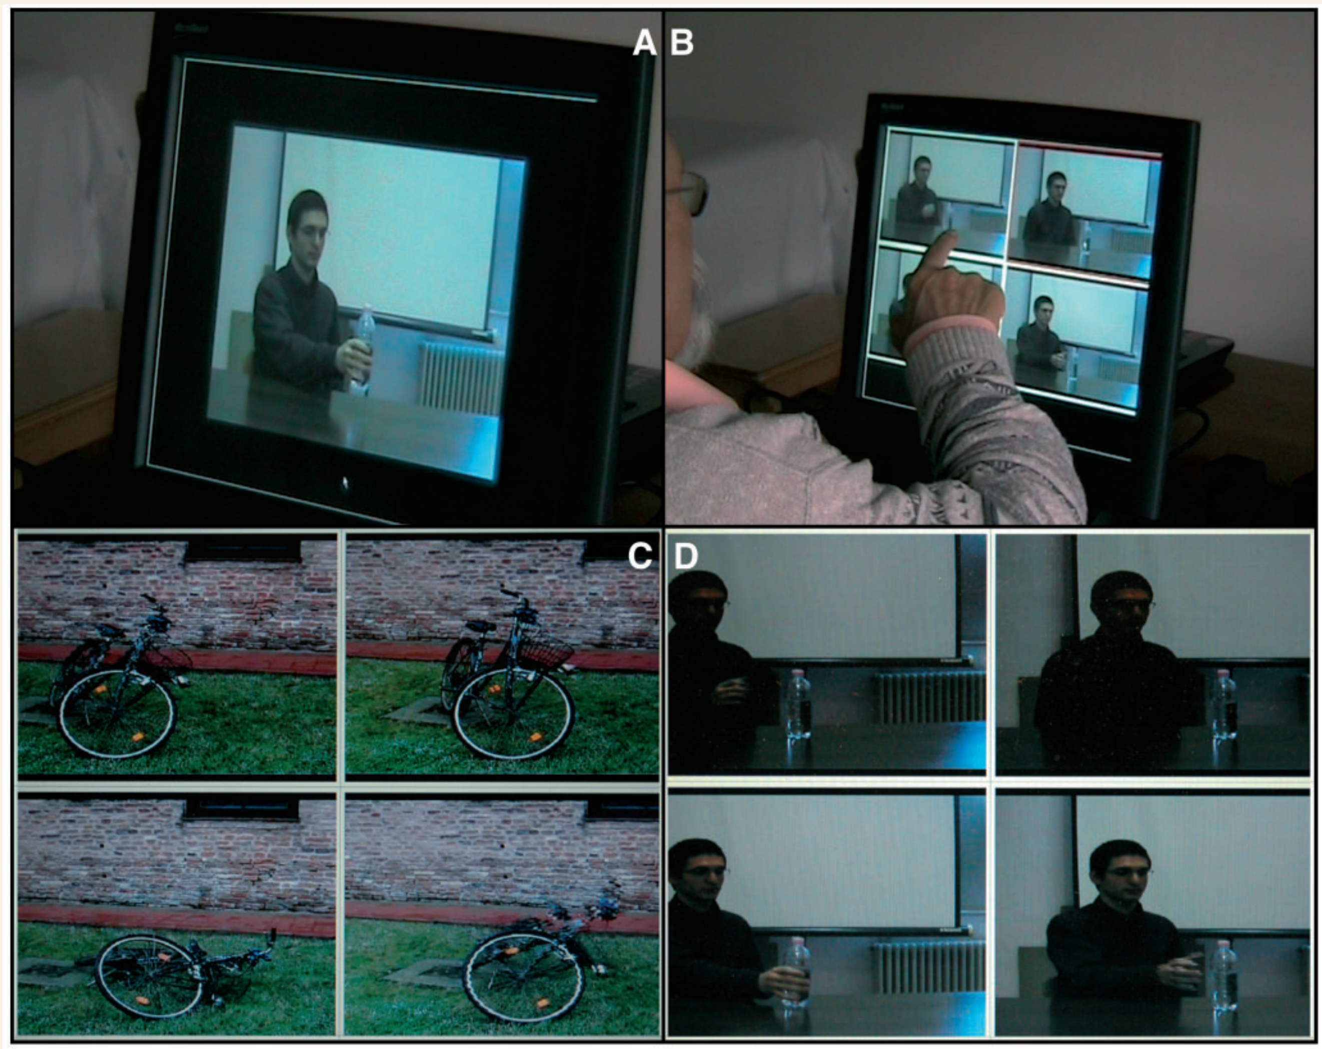

Figure 2 Experimental set-up and task. (**A**) The videoclip is presented on the screen, (**B**) Four snapshots are presented at the four corners of the screen, (**C**) Example of physical events snapshots and (**D**) Example of Human Actions snapshots.

sequencing task and in the language sequencing task (LST). This analysis served to measure whether there was a performance correlation across the two tests. All the analyses were performed using Statistica 6 (StatSoft, Inc.).

Results

The score in the neuropsychological test battery was on average 23.5 out of 29 items on the apraxia section with a SD of 2.88. The controls obtained 28.5/29 and a SD of 0.55. The language section showed a performance of 43.5 out of 67 with a SD of 15.1 for patients, while controls scored on average 64.67 out of 67 items and a SD of 1.63. Among all neuropsychological tests, the denomination of tools, actions and tool use showed a significantly worse performance in patients ($U = 32.5$, $P < 0.05$, $U = 36$, $P < 0.01$ and $U = 25.5$, $P = 0.24$, respectively), whereas denomination of natural objects was not impaired ($U = 25.5$, $P = 0.24$). Table 3 shows the performance of patients and controls on these items of the neuropsychological questionnaire.

The ANOVA on RTs [main effect: CONDITION, $F(1,5) = 1.675$; $P = 0.25$. GROUP: $F(1,5) = 15.95$; $P = 0.01$; Interaction CONDITION $\times$ GROUP, $F(1,5) = 0.026$; $P = 0.88$] as well as the ANOVA on TT [main effect: CONDITION, $F(1,5) = 0.406$; $P = 0.55$. GROUP: $F(1,5) = 24.665$; $P = 0.004$; Interaction CONDITION $\times$ GROUP, $F(1,5) = 0.439$; $P = 0.54$] showed a significant effect for factor GROUP, indicating that patients were generally slower than controls, independent of the experimental manipulation (physical events or human actions). The speed of performance (RT and TT) was not specifically influenced by the experimental condition (human actions and physical events), except for the fact that the healthy controls were consistently faster than the patients. This speed bias can easily be accounted for by the general increase in reaction times and in movement times often observed in brain-lesioned patients (Benton, 1986).

More interestingly, ANOVA performed on Accuracy showed a significant effect for the interaction GROUP $\times$ CONDITION

Table 2 Stimuli list

Stimuli	RT	TM	TT
Touching the tip of one's nose	2.133 ± 0.724	7.058 ± 5.125	9.191 ± 5.017
A bow	2.317 ± 1.354	4.683 ± 2.46	7 ± 2.9
Climbing a ladder to get a box	2.358 ± 0.643	7.767 ± 3.515	10.125 ± 3.908
A bicycle falling*	2.459 ± 0.525	4.017 ± 0.653	6.475 ± 0.918
Approaching a wall on all fours and touching it	2.484 ± 0.759	5.725 ± 1.75	8.208 ± 1.436
Plotter*	2.691 ± 1.21	4.267 ± 2.331	6.958 ± 2.536
Grabbing a bottle	2.767 ± 1.226	7.208 ± 3.008	9.975 ± 4.065
Turning one's head and pointing	2.991 ± 1.431	5.467 ± 3.307	8.458 ± 3.413
Cutting a sheet of paper with a pair of scissors	3.208 ± 1.291	6.075 ± 1.872	9.283 ± 2.081
Opening a wardrobe by turning the key	3.392 ± 1.572	10.392 ± 7.347	13.783 ± 6.739
Opening a notebook and writing	3.508 ± 1.185	10.975 ± 6.818	14.483 ± 7.274
A remote controlled car against a wall*	3.667 ± 1.593	6.542 ± 1.52	10.208 ± 1.723
Getting over an obstacle	3.725 ± 1.346	6.258 ± 3.058	9.983 ± 3.128
Getting up from the ground	3.783 ± 2.137	7 ± 2.607	10.783 ± 4.322
A door closing*	3.95 ± 2.457	5.692 ± 1.69	9.642 ± 3.458
Taking off one's glasses	3.975 ± 1.47	7.45 ± 5.832	11.425 ± 6.943
A ball rolling down an inclined plane*	4.1 ± 1.762	6.517 ± 2.203	10.617 ± 2.823
Wiping out a blackboard	4.133 ± 1.967	8.7 ± 5.328	12.833 ± 6.041
Opening a wallet and take out an ID	4.158 ± 2.233	5.933 ± 2.753	10.092 ± 4.455

List of 19 movies presented to both patients and matched controls. RT, TM, TT and relative SD for each stimulus, measured during the pilot experiment in 13 subjects, are provided. Asterisks denote physical events stimuli.

Table 3 Neuropsychological testing performance for action-related items

Denomination	Objects	Tools	Tool uses	Action
Patients				
FG	3	1	0	0
GF	11	5	3	3
SC	12	7	4	6
CC	9	1	2	0
DF	10	4	0	2
EC	12	5	6	6
Mean	9.5/12	3.83/8	2.5/8	2.83/8
Control				
AF	12	8	8	8
GB	12	7	8	8
VG	11	6	8	8
VV	12	8	8	8
ES	11	6	8	8
MT	11	6	8	8
Mean	11.5/12	6.83/8	8/8	8/8

Patient's and matched control's performance on denomination of objects, actions, tools and tools use items present in the neuropsychological testing. Performance between the two groups did differ significantly in all items but not in the denomination of natural objects.

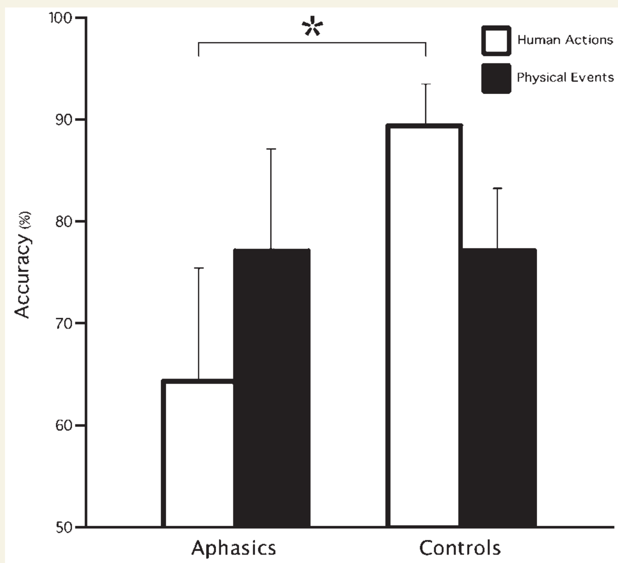

Figure 3 Accuracy results. Histograms depict the accuracy ratio in aphasic patients (Aphasics) and normal subjects (Controls) for both human actions (white bars) and physical events (black bars) conditions. Whiskers indicate the standard error of the mean. Asterisks denote statistically significant differences ($P < 0.05$) in accuracy ratio between aphasics and controls in the human action condition.

[$F(1,5) = 12.594$; $P = 0.02$], but no significant effect for factor GROUP [$F(1,5) = 4.314$; $P = 0.09$] and CONDITION [$F(1,5) = 0.0005$; $P = 0.98$]. *Post hoc* comparisons showed that while the control group performed equally well in the two experimental conditions (correct response for human actions: 0.89 ± 0.04; correct response for physical events 0.77 ± 0.06 SEM; $P =$ NS), aphasics' performance showed a trend to significance ($P = 0.05$) between human actions (0.64 ± 0.11) compared with physical events (0.77 ± 0.10), and a highly significant difference with respect to controls for the human action condition ($P = 0.004$) (Fig. 3).

A further analysis was carried out on the intransitive versus transitive human action trials. To this purpose, we separated patients' and controls' data into transitive ($n = 8$) and intransitive trials ($n = 6$). As transitive video clips we defined hand–object interactions (e.g. grasping a glass). The intransitive video clips were those representing meaningful actions but without object (e.g. turning the head and point). Paired *t*-tests showed that patients were significantly impaired with respect to controls in both transitive ($P = 0.04$) and intransitive trials ($P = 0.001$), without any difference in performance between transitive and intransitive trials ($P = 0.37$). Therefore, results cannot be due to the presence of a human–object interaction, since a similar impairment was also present for the intransitive human actions.

The language sequencing test showed a severe impairment for patients in the verbal domain and low accuracy levels in reorganizing scrambled sentences or syllables (LST, mean percentage of correct responses $\pm$ SEM: All LST: 53.15 ± 12.24; sentences: 52.08 ± 15.95; syllables: 53.64 ± 12.41), confirming their deficits in the language domain. Moreover the correlation analysis between the action- and language-sequencing test showed a significant relation between performance in sequencing transitive, human action video clips and all items of the language sequencing test ($r^2 = 0.74$; two-tailed $P = 0.03$). Conversely, performance in the sequencing of intransitive actions and physical events was not correlated with the language sequencing test (Intransitive actions: $r^2 = 0.04$; two-tailed $P = \mathrm{NS}$; physical events: $r^2 = 0.15$; two-tailed $P = \mathrm{NS}$).

Discussion

The present work shows that frontal aphasic patients, characterized by a lesion centred in the left *pars opercularis* of Broca's region and by the absence of apraxic symptoms, are specifically impaired in sequencing pictures representing actions (transitive or intransitive) performed by a human agent but not in sequencing physical events. Additionally, their reduced ability to sequence sentence segments and word syllables correlated with the impairment in sequencing transitive actions. Although, it is still possible that plastic processes and/or compensatory strategies might take advantage from the right homologue region, patients did not restore these specific abilities.

Why are these patients not able to solve the sequencing task for human actions only? In our experiment, subjects were requested to understand what they were observing in the video clip, and then order single snapshots into a meaningful sequence. To do this, we suppose the subjects had to represent (and replay) the rules connecting critical information presented in the videos. The interpretation we favour is that, to correctly sequence human actions, subjects were implicitly mapping the observed actions onto their own motor repertoire. In other words, the subjects had to gain access to 'how' a given action was composed in terms of simple units, and harmonically (and pragmatically) restructure it through an embodiment process. Conversely, in the case of physical events, such an implicit and embodied motor representation was unnecessary to solve the task. This interpretation is in line with the finding that similar results were achieved in sequencing both transitive and intransitive actions, and complement the idea that the human-object interaction is not a necessary prerequisite to activate the motor system during action observation (Fadiga *et al*., 1995). More interestingly, patients' performance in human action sequencing was also correlated to their deficit in sequencing words forming sentences and syllables forming words. Moreover, they all had severe problems in naming tools and, more importantly, tools' uses. On the contrary, patients' understanding of the global meaning of the observed actions was mostly preserved if they were asked to explain what they had seen.

Why should this capacity of representing action pragmatics be encoded in Broca's area? A large number of recent neuroimaging and neurophysiological studies have shown that a reproducible network of cortical areas, comprising Broca's region, becomes active during action observation (for a review see Rizzolatti and Craighero, 2004; for a critical position on the possibility to consider these activations as a proof of the existence of a human mirror-neuron system see Turella *et al*., 2008).

This productive area of research has been motivated by previous monkey studies describing similar mechanisms at a cellular level in macaque premotor area F5 (di Pellegrino *et al*., 1992; Gallese *et al*., 1996; Rizzolatti *et al*. 1996) and in the inferior parietal lobule (Fogassi *et al*., 2005; Rozzi *et al*., 2008).

Frontal and parietal mirror neurons found in the macaque brain fire when the monkey executes an action and also when it observes the same action performed by someone else. It has been suggested that mirror neurons may provide the brain with an implicit knowledge about the meaning of actions because seen actions are directly matched onto the observer's motor repertoire. Therefore, the finding that Broca's area, the putative human cytoarchitectonic homologue to monkey area F5 (Petrides and Pandya, 1997; Petrides *et al*., 2005), becomes active during action observation, strongly supports the hypothesis that it may form a crucial node of the human mirror-neuron system. It could be a wrong, or at least too simplistic a conclusion, to think that Broca's area and its monkey homologue share all their functional properties. Indeed, evolution is characterized by an increase of cytoarchitectonically diverse cortical areas. For this reason, the functional properties of monkey area F5 might have been distributed to different sectors of the human premotor cortex, probably according to their degree of response complexity. However, functional and anatomical evidence reinforce the intriguing possibility that the goal-related action vocabulary stored in monkey premotor cortex (Rizzolatti *et al*., 1988) and the syntax-related properties of Broca's area, might be evolutionarily linked. In our view the data presented by our work strengthens this link by providing, for the first time, clear evidence that Broca's aphasics show a significant impairment in representing observed actions.

Actions, by definition, are hierarchical compositions of simpler motor acts (Grafton and Hamilton, 2007) aiming at a goal. Thus, action decoding via visual information may require the harmonic composition of low level visual-kinematic features into a high level representation of action-goals and therefore of agent's intention. The same intention can be conveyed by a set of movements with quite a large degree of inter- and intra-subject variability, which the system has to efficiently categorize as pertaining to the same action. What remains constant and so critically useful, are the rules to compose such hierarchically lower units. We might consider this set of rules as a sort of motor syntax, the knowledge of which is, in our view, necessary to solve our sequencing task in the case of human actions but not in that of physical events.

In agreement with our interpretation, Dominey *et al*. (2003) and Sirigu *et al*. (1998) demonstrated that patients with lesions of Broca's area are impaired in learning the hierarchical/syntactic structure of linguistic sequential tasks. Moreover, and more recently, an event-related fMRI study succeeded in disentangling hierarchical processes from temporally nested elements (Koechlin and Jubault, 2006). These authors reported that Broca's area, and its right homologue, control selection and nesting of action segments, integrated in hierarchical behavioural plans, regardless of their temporal structure. Finally, Bahlmann *et al*. (2008) showed that, when comparing the processing of hierarchical dependencies to adjacent dependencies in an artificial language, significantly higher activations were observed in Broca's area and in the ventral premotor cortex. These results indicate that Broca's area may form a node of a neural circuit responsible for processing hierarchical structures in an artificial grammar context.

In our view, Broca's area might have specialized in encoding complex hierarchical structures of goal-directed actions, and to eventually apply these pragmatic rules to more abstract domains. Therefore, the language-related functions sub-served by Broca's region could be the most eloquent part of a more general computational mechanism shared by multiple domains. Such mechanisms could be imagined as a polymodal syntax (Baumgaertner *et al*., 2007) endowed with the ability to organize and comprehend hierarchically dependent elements into meaningful verbal and non-verbal structures.

Conclusions

The present work sheds light on the functional role of Broca's area by providing evidence that, in the absence of apraxia, a lesion affecting Broca's area impairs the ability to sequence actions in a task with no explicit linguistic requirements. Here, we propose that the complex pattern of abilities associated with Broca's area might have evolved from its premotor function of assembling individual motor acts into goal-directed actions. This capacity of dealing with complex motor hierarchical structures could have evolved into a polymodal syntax serving also higher cognitive functions sharing with action some basic grammatical rules. Consequently, we speculate that an ancient motor syntax might have evolved into a 'supramodal syntax', at the basis of the 'modern' linguistic one.

Funding

European Community grants (CONTACT-NEST Project 5010, ROBOTCUB-IST-004370, POETICON-ICT-215843); Italian Ministry of Education grants (to L.F.).

References

Alexander MP, Naeser MA, Palumbo C. Broca's area aphasia: aphasia after lesions including the frontal operculum. Neurology 1990; 40: 353–62.

Amunts K, Schleicher A, Bürgel U, Mohlberg H, Uylings HB, Zilles K. Broca's region revisited: cytoarchitecture and intersubject variability. J Comp Neurol 1999; 412: 319–41.

Bahlmann J, Schubotz RI, Friederici AD. Hierarchical artificial grammar processing engages Broca's area. Neuroimage 2008; 42: 525–34.

Baumgaertner A, Buccino G, Lange R, McNamara A, Binkofski F. Polymodal conceptual processing of human biological actions in the left inferior frontal lobe. Eur J Neurosci 2007; 25: 881–9.

Bell BD. Pantomime recognition impairment in aphasia: an analysis of error types. Brain Lang 1994; 47: 269–78.

Benton A. Reaction times in brain disease: some reflections. Cortex 1986; 22: 129–40.

Binkofski F, Buccino G, Posse S, Seitz RJ, Rizzolatti G, Freund H. A fronto-parietal circuit for object manipulation in man: evidence from an fMRI-study. Eur J Neurosci 1999; 11: 3276–86.

Broca MP. Remarques sur le siege de la facultè du langage articulé, suivies d'une obserevation d'aphemie (Perte la parole). Bulletins et memoires de la Societe Anatomique de Paris 1861; 36: 330–57.

Buccino G, Binkofski F, Fink GR, Fadiga L, Fogassi L, Gallese V, Seitz RJ, Freund HJ. Action observation activates premotor and parietal areas in a somatotopic manner: an fMRI study. Eur J Neurosci 2001; 13: 400–4.

Capasso R. Esame neuropsicologico per l'Afasia (E.N.P.A.). Milano: Springer-Verlag Italia; 2001.

Caplan D, Hildebrandt N, Makris N. Location of lesions in stroke patients with deficits in syntactic processing in sentence comprehension. Brain 1996; 119: 933–94.

Caramazza A, Zurif EB. Dissociation of algorithmic and heuristic processes in language comprehension: evidence from aphasia. Brain Lang 1976; 3: 572–82.

Ciurli P, Marangolo P, Basso A. Esame del linguaggio—II. Firenze: Giunti Organizzazioni Speciali 1996; 00: 0.

Daniloff JK, Noll JD, Fristoe M, Lloyd LL. Gesture recognition in patients with aphasia. J Speech Hear Disord 1982; 47: 43–9.

De Renzi E, Motti F, Nichelli P. Imitating gestures: a quantitative approach to ideomotor apraxic. Arch Neurol-Chicago 1980; 37: 6–10.

De Renzi E, Pieczuro A, Vignolo LA. Oral apraxia and aphasia. Cortex 1966; 2: 50–73.

De Renzi E, Vignolo LA. The token test: a sensitive test to detect receptive disturbances in aphasics. Brain 1962; 85: 665–78.

Di Pellegrino G, Fadiga L, Fogassi L, Gallese V, Rizzolatti G. Understanding motor events: a neurophysiological study. Exp Brain Res 1992; 91: 176–80.

Dinstein I, Thomas C, Behrmann M, Heeger DJ. A mirror up to nature. Curr Biol 2008; 18: R13–8.

Dominey PF, Hoen M, Blanc JM, Lekelov-Boissard T. Neurological basis of language and sequential cognition: evidence from simulation, aphasia, and ERP studies. Brain Lang 2003; 86: 207–25.

Dronkers NF. A new brain region for coordinating speech articulation. Nature 1996; 384: 159–61.

Dronkers NF, Plaisant O, Iba-Zizen MT, Cabanis EA. Paul Broca's historic cases: high resolution MR imaging of the brains of Leborgne and Lelong. Brain 2007; 130: 1432–41.

Duffy RJ, Duffy JR. Pantomime recognition in aphasics. J Speech Hear Res 1975; 18: 115–32.

Evans C, Collins DL, Mills SR, Brown ED, Kelly RL, Peters TM. 3D statistical neuroanatomical models from 305 MRI volumes. Proceedings of the IEEE-nuclear science symposium and medical imaging conference. 1813–7.

Fadiga L, Craighero L, Roy AC. Broca. In: Grodzinsky Y, Amunts K, editors. Broca's region. New York: Oxford University Press; 2006. p. 137–52.

Fadiga L, Fogassi L, Pavesi G, Rizzolatti G. Motor facilitation during action observation: a magnetic stimulation study. J Neurophysiol 1995; 73: 2608–11.

Fogassi L, Ferrari PF, Gesierich B, Rozzi S, Chersi F, Rizzolatti G. Parietal lobe: from action organization to intention understanding. Science 2005; 308: 662–7.

Gainotti G, Lemmo MS. Comprehension of symbolic gestures in aphasia. Brain Lang 1976; 3: 451–60.

Gallese V, Fadiga L, Fogassi L, Rizzolatti G. Action recognition in the premotor cortex. Brain 1996; 119: 593–609.

Glosser G, Wiener M, Kaplan E. Communicative gestures in aphasia. Brain Lang 1986; 27: 345–59.

Goldenberg G. Defective imitation of gestures in patients with damage in the left or right hemispheres. J Neurol Neurosur Psychiatry 1996; 61: 176–80.

Grafton ST, Arbib MA, Fadiga L, Rizzolatti G. Localization of grasp representations in humans by PET: 2. Observation compared with imagination. Exp Brain Res 1996; 112: 103–11.

Grafton ST, Hamilton AF. Evidence for a distributed hierarchy of action representation in the brain. Hum Mov Sci 2007; 26: 590–616.

Grèzes J, Decety J. Functional anatomy of execution, mental simulation, observation, and verb generation of actions: a meta analysis. Hum Brain Map 2001; 12: 1–19.

Iacoboni M, Woods R, Brass M, Bekkering H, Mazziotta JC, Rizzolatti G. Cortical mechanisms of human imitation. Science 1999; 286: 2526–8.

Koechlin E, Jubalt T. Broca's area and the hierarchical organization of human behaviour. Neuron 2006; 50: 963–74.

Murata A, Fadiga L, Fogassi L, Gallese V, Raos V, Rizzolatti G. Object representation in the ventral premotor cortex (area F5) of the monkey. J Neurophysiol 1997; 78: 2226–30.

Nishitani N, Hari R. Temporal dynamics of cortical representation for action. Proc Natl Acad Sci USA 2000; 97: 913–8.

Ojemann G, Ojemann J, Lettich E, Berger M. Cortical language localization in left, dominant hemisphere. An electrical stimulation mapping investigation in 117 patients. J Neurosurg 1989; 71: 316–26.

Pazzaglia M, Smania N, Corato E, Aglioti SM. Neural underpinnings of gesture discrimination in patients with limb apraxia. J Neurosci 2008; 28: 3030–41.

Penfield W, Roberts L. Speech and brain mechanisms. Princeton NJ: Princeton University Press; 1959.

Petrides M, Cadoret G, Mackely S. Oro-facial somatomotor response in the macaque monkey homologue of Broca's area. Nature 2005; 435: 1235–8.

Petrides M, Pandya DN. Comparative architectonic analysis of the human and the macaque frontal cortex. In: Boller F, Grafman J, editors. Handbook of neuropsychology. New York: Elsevier; 1997. p. 17–58.

Pobric G, Hamilton AF. Action understanding requires the left inferior frontal cortex. Curr Biol 2006; 16: 524–9.

Rizzolatti G, Camarda R, Fogassi L, Gentilucci M. Functional organization of inferior area 6 in the macaque monkey II. Area F5 and the control of distal movements. Exp Brain Res 1988; 71: 491–507.

Rizzolatti G, Craighero L. The Mirror neuron system. Ann Rev Neurosci 2004; 27: 169–92.

Rizzolatti G, Fadiga L, Gallese V, Fogassi L. Premotor cortex and the recognition of motor actions. Cog Brain Res 1996; 3: 131–41.

Rorden C, Brett M. Stereotaxic display of brain lesions. Behav Neurol 2000; 12: 191–200.

Rozzi S, Ferrari PF, Bonini L, Rizzolatti G, Fogassi L. Functional organization of inferior parietal lobule convexity in the macaque monkey: electrophysiological characterization of motor, sensory and mirror responses and their correlation with cytoarchitectonic areas. Eur J Neurosci 2008; 28: 1569–88.

Saygin AP, Wilson MS, Dronkers NF, Bates E. Action comprehension in aphasia: linguistic and non-linguistic deficits and their lesion correlates. Neuropsychologia 2004; 42: 1788–804.

Schäffler L, Lueders HO, Beck GJ. Quantitative comparison of language deficits produced by extraoperative electrical stimulation of Broca's, Wernicke's, and Basal temporal language areas. Epilepsia 1996; 37: 463–75.

Schäffler L, Lüders HO, Dinner DS, Lesser RP, Chelune GJ. Comprehension deficits elicited by electrical stimulation of Broca's area. Brain 1993; 116: 695–715.

Sirigu A, Cohen L, Zalla T, Pradat-Diehl P, Van Eeckhout P, Grafman J, Agid Y. Distinct frontal regions for processing sentence syntax and story grammar. Cortex 1998; 34: 771–8.

Talairach J, Tournoux P. Co-planar stereotaxic atlas of the human brain: 3-dimensional proportional system - an approach to cerebral imaging. New York: Thieme Medical Publishers]; 1988.

Tranel D, Kemmerer D, Adolphs R, Damásio H, Damásio AR. Neural correlate of conceptual knowledge for actions. Cogn Neuropsychol 2003; 20: 409–32.

Turella L, Pierno AC, Tubaldi F, Castiello U. Mirror neurons in humans: consisting or confounding evidence? Brain Lang 2009; 108: 10–21.

Varney NR. Pantomime recognition defect in aphasia: implication for the concept of asymbolia. Brain Lang 1982; 15: 32–9.

Wang L, Goodglass H. Pantomime, praxis, and aphasia. Brain Lang 1992; 42: 402–18.

PROSOCIAL ATTITUDE MODULATES MUSCLE ACTIVITY IN A SIMPLE COORDINATION GAME

Roberto Censolo[1#], Laila Craighero[1], Giovanni Ponti[1], Leonzio Rizzo[1], Luciano Fadiga[1, 2#]

[1] University of Ferrara, Ferrara, Italy

[2] The Italian Institute of Technology, Genova, Italy

Corresponding author:
Prof. Luciano Fadiga, Section of Human Physiology, University of Ferrara, 44100 Ferrara, Italy
Phone: +39-0532-455241 Fax: +39-0532-455242

ABSTRACT

Social and economic literature generally considers that the relevance of property rights in human interaction arises from an explicit cognitive processes, which emerge with social competence. In the present work we aimed at investigating whether the arbitrary allocation of property rights automatically affects individuals' behavior, within a context with no visual and verbal interaction and removed from any perspective-taking activity. To this purpose we submitted twelve pairs of participants to a simple motor coordination game while recording the electromyographic activity of the muscle mainly involved in the task. By using the answers to a questionnaire to measure the participants' degree of sociability, the correlation between muscle involvement and prosocial attitude revealed that low prosocial individuals, only, significantly changed their motor behavior in response to a reallocation of property rights. Results are discussed in relation to the *endowment effect*, a puzzling phenomenon observed in actual behavior, that had challenged the traditional assumption of rationality in many behavioral models.

INTRODUCTION

In actual societies the most part of collective welfare is generated through market transactions and social services. This bulk of social coordination requires a common set of shared norms to effectively take place, essentially property rights on things and actions. The main focus of the present paper was to investigate if the effectiveness of property rights should be entirely ascribed to cognitive processes related to strategic or "perspective-taking" considerations arising with social competence, or whether they constitute a set of behavioral devices able to automatically influence individual's behavior.

The role of property rights in social interaction have been extensively investigated within the experimental economics framework. A conspicuous amount of recent research has focused attention to investigate the influence of antecedents on the perceived payoffs of others in strategic environments. Typically, the outcomes of a game under perfect anonymity are compared to those obtained in a two stages experimental design. In the first stage one (or both) players accomplish a task by virtue of which they acquire a "role". In the second stage, subjects interact strategically, knowing nothing of each other except for the role gained in the first stage. Overall, empirical findings support the view that the subjects' behavior, and the distribution of payoffs, considerably reflect the "entitlements" earned by participants. For example, in the 'dictator game' the first player (proposer) has to divide a sum between herself and the second player (receiver), who passively receives the share allocated to her. In this context, if the proposer earned some entitlement[1] to the sum assigned to her, the frequency of zero offers[2] to the receiver increases, while positive offers arise more frequently on the part of the proposer if the receiver gained some "role" (Hoffman et al. ;1992, Oxoby and Spraggon; 2008, Cherry; 2002). The 'ultimatum game' is a strategic version of the dictator, widely employed in this literature, where the receiver can either accept or refuse the offer. If the receiver accepts, the stake is split according to the proposal. If she rejects, neither player get anything. In this setting, if the receiver earned a legitimate role, she gets larger shares of money from the first mover (Ruffle; 1998, Cherry; 2001). Overall, this literature

[1] Entitlement may be gained by scoring high on a general knowledge quiz (Hoffman et al.; 1992; Hoffman et al.; 1996, Cherry; 2002) or even by cracking a sufficient amount of walnuts (Fahr and Irlenbush; 2000).
[2] A zero offer corresponds to the standard selfishness-rationality prediction.

shows that having obtained a role by the accomplishment of a specific task, subjects appear to consider that one has a right to outcomes, which in other circumstances may be regarded as unjust. This body of evidence has led to support the interpretation that the perception of *legitimate property rights* on the part of individuals constitutes an important element influencing social interaction.

However, the experimental designs by which these results are obtained, make it difficult to asses the relevance of the subjective perception of property rights clearly distinguished from the legitimation sources of such entitlements. Indeed, the acquisition of a role seems to condition actual behavior of participants through the modulation of other emotional and/or strategic aspects relevant to the decision process. For example, in the dictator game the asymmetry in the perception of the "other" introduced by the distribution of "roles" might affect the participants' sense of equity, or, in the ultimatum context, it might influence the strategic assessment of the risk of rejection. In this respect, the experimental setup designed by this research neglects the possibility that the subjective perception of "property rights" might represent a distinct dimension along which social interaction takes place.

With respect to this approach, the main focus of our experiment was to verify whether the formal entitlements of property rights, regardless of any legitimating activity undertaken by participants, play a significant role within a context where interaction between individuals does not involve any explicit process related to emotional cues and/or to strategic or "perspective-taking" considerations. Furthermore, we wanted to investigate if the allocation of property right automatically influence individual's behavior at a very low level, such as the intensity of muscle involvement during the execution of hand actions. To this purpose, twelve pairs of participants, prevented from any visual or verbal exchange, were submitted to a simple motor coordination task. Each couple had to cooperatively hold a small sphere between their right index fingers and to drop it alternately into one of two containers placed below their hands, while electromyography of the right *first dorsal interosseus* (FDI) muscle of each participant was recorded. Each successful trial was differently rewarded with a given amount of money according to the experimental condition, and the rewarding rules were communicated before starting each session. Consequently, for the same action (e.g., pushing the ball into the leftside container) each participant could receive a reward in one session but not in another. The total monetary reward gained by each subject in each condition was always the same. Finally, we correlated muscle involvement to the scores obtained in a social attitude questionnaire to verify if the degree of prosocial propensity covertly modulates motor behavior.

METHODS

Subjects. Twenty-four female participants were recruited among students of the Law Department of the University of Ferrara (mean age 26 +/- 3). All of them were naïve to the purpose of the experiment, were right-handed according to the Oldfield questionnaire (Oldfield, 1971) and gave their informed consent. They were divided into two subgroups (the "Green" and the "Yellow" group) of 12 participants, and kept in separate rooms after their arrival at the lab. Twelve pairs of subjects were then formed by extracting randomly one partner from each subgroup. Each pair, composed by one Green and one Yellow subject, was submitted to an experimental session lasting approximately 30 minutes.

Questionnaire. In the first stage of the experiment the subjects were asked to answer a written questionnaire based on Putnam's *Social Capital Benchmark Survey* (http://www.hks.harvard.edu/saguaro/communitysurvey/index.html). Following Bobo et al. (1995) we employed the answers provided by subjects to build several indexes aimed at measuring individual prosocial/proself attitude (see the Appendix for details).

Coordination game Before entering the lab room, subjects have been invited to remove rings, bracelets, nail enamel, or other kind of decoration, that could have made them recognizable by other subjects. At the beginning of the experiment, two subjects entered the experimental room from two different doors, standing one in front of the other, their face and trunk hidden by a curtain. Thus, during the experimental session subjects never saw each other. Moreover, they were strictly recommended not to speak to exclude any possible recognition based on subject's voice.

Subjects were requested to pose their forearms on a Plexiglas surface with a square hole in correspondence of their hands. Below the Plexiglas was set an apparatus constituted by two adjacent containers of equal size, with the partition side aligned with participants' sagittal plane. At the beginning of each trial a small glass sphere (1 cm diameter) was placed between the extended right index fingers of the two subjects, and subjects were requested to stay on this position (starting position) until the go-signal. In this position the sphere was exactly above the border between the two containers. Subjects' index fingers were dressed with a soft sponge to avoid flexion in the course of the play, and to increase the attrition surface to better keep the sphere in the proper position.

Each pair of subjects was asked to play 30 trials of a simple motor ability game. The 30 trials were subdivided into three experimental conditions (C_1, C_2 and C_3) of ten trials each, blocked into three experimental sessions, the presentation of which was pseudo-randomly balanced across pairs. At every trial subjects followed the instruction given by the experimenter indicating to drop the sphere alternately into the two containers. The difference among conditions C_1, C_2 and C_3. consisted in the monetary incentive associated to each trial successfully performed by subjects. Specifically, in condition 1, putting the sphere into either target container yielded a reward of € 0.50 to each subject (Figure 1A). In Condition 2 and 3, two colored sheets, one green and one yellow, were placed onto the floor of each container, defining the Green and the Yellow container. The allocation of rewards coupled containers and subjects of the same color. When the sphere was successfully dropped into the target container a € 1 reward was received by the correspondent colored subject only. In Condition 2, the Green (Yellow) container was placed at the left side of the Green (Yellow) subject: the winning subject had to execute an index finger abduction (contraction of the FDI muscle) to push the sphere towards the container (Figure 1B). In Condition 3, the colors of containers were reversed, so that the Green (Yellow) container was placed at the right side of the Green (Yellow) subject: the winning subject had to execute an index finger adduction (FDI muscle not involved) to "pull" the sphere towards the container (Figure 1C). The total money reward gained by each subject was € 5 in each condition (€ 15 total).

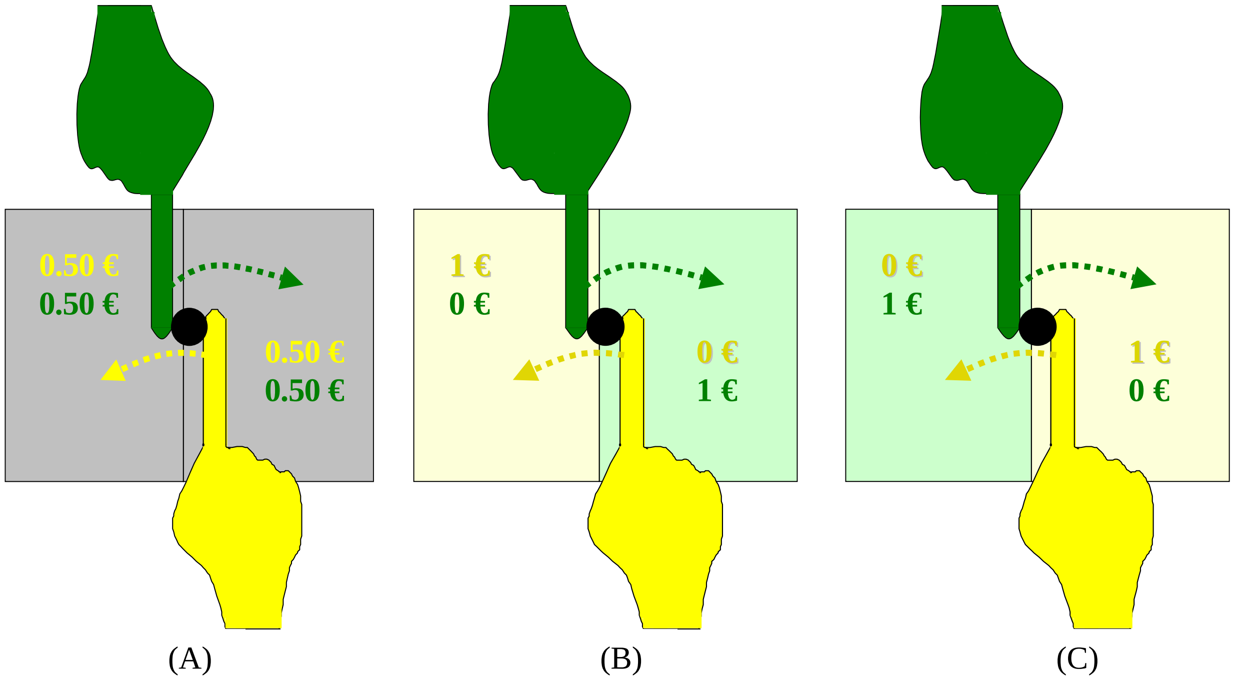

Figure 1. Schematic view of the experimental apparatus used in the three experimental conditions. Subjects' hands laid on a Plexiglas plate with the two index fingers positioned in correspondence of a square hole (the rectangle shown in the figure). Under the Plexiglas plate, at a distance of 10 cm from it, there were two containers (the two grey areas shown in panel A) where the subjects had to drop the sphere held by their index fingers according to the specific instructions provided for each experimental conditions. The moment at which the sphere touched the floor of the container was detected by a load cell**.** The monetary incentives associated to the three experimental conditions were the following: Condition 1 (A): each subject (Yellow and Green) get € 0.50 at any trial. Condition 2: the Yellow (Green) subject is coupled with the Yellow (Green) container; the pushing subject gets € 1 while the pulling one gets zero. Condition 3: the container are reversed; the pushing subject gets zero and the pulling one gets € 1. Ten trial for each condition. Each subject received € 5x3 = € 15.

Electromyographic potentials (EMG) were recorded from right *first dorsal interosseus* (FDI) muscle by using Ag-AgCl surface electrodes (diameter 6 mm) glued to the subjects' skin according to a tendon-belly configuration. After online rectification and integration (time constant 0.05 s) EMG signal was continuously recorded during the experiment and fed to a personal computer for the successive analysis. The acquisition software sampled the EMG signal recorded from the two subjects at 25 Hz. The instant at which the ball touched the bottom of the target container was detected by means of a load cell supporting the container itself. The load cell signal, appropriately amplified, was continuously acquired during the experiment by the same acquisition software used for EMG recordings (at the same sampling frequency).

DATA ANALYSIS AND RESULTS

As Table 1 shows, subjects were able to coordinate almost perfectly in all three experimental conditions, with only a negligible proportion of inefficient outcomes (2.7% of total observations), uniformly distributed across conditions.

Condition	*Green wins*	*Yellow wins*	*Inefficient outcomes*	*Total (12 pairs x 10 trials)*
1	58	58	4	120
2	59	58	3	120
3	59	58	3	120
Total	176	174	10	360

Table 1. Outcomes of the game for each condition

Figure 2 depicts the typical EMG traces recorded from both subjects' FDI muscles (blue and red traces) and the signal recorded from the load cell, detecting the instant at which the sphere, after its releasing, touches the floor of the container (black trace), during condition 1 (A) and 3 (B).

(A)

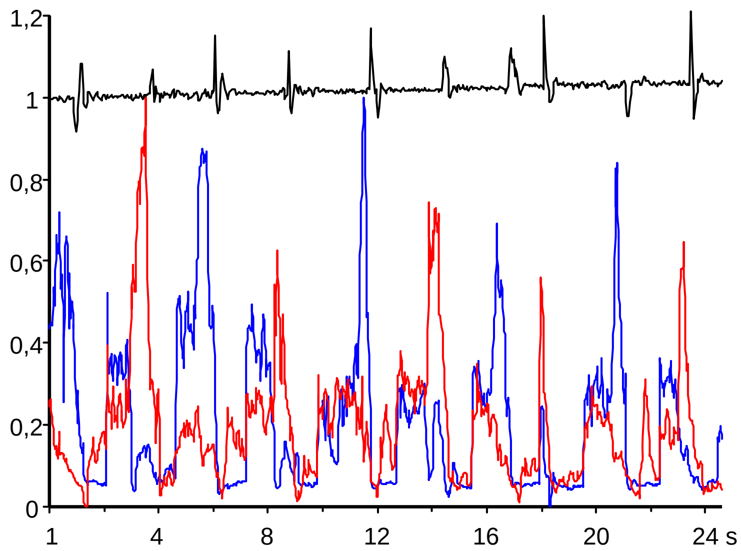

(B)

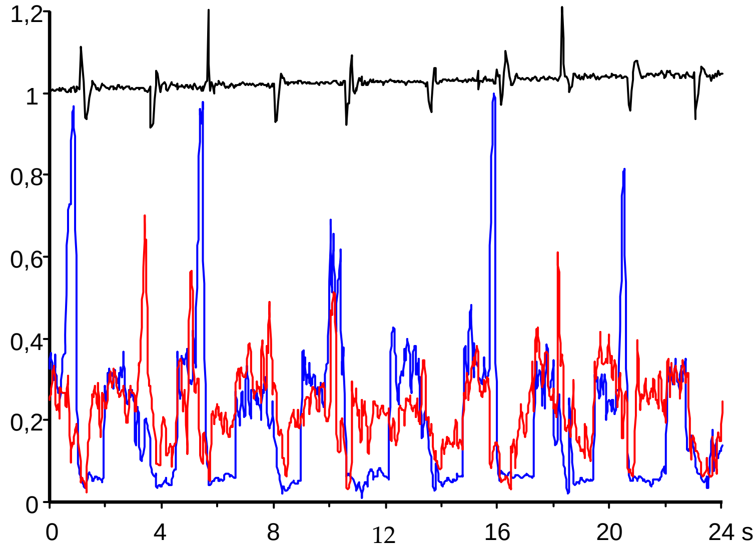

Figure 2. Typical *first dorsal interosseus* electromyography rectified, integrated (time constant, 0.05 s) and intra-subject normalized (z-scores), as recorded from two subjects (red and blue traces) during the interaction game. Panel A, Condition 1; panel B, Condition 3. The signal recorded from the load cell is shown in black and indicates the ten times the glass sphere fell into the container, signaling the end of each trial. The figure depicts ten subsequent trials (sampling frequency, 25 Hz). Abscissas, seconds; ordinates, arbitrary normalization units (see text).

As it appears from Figure 2, at the beginning of each trial there is an increase of both subjects' EMG determined by the involvement of subjects' index fingers in maintaining the glass sphere in the starting position. After the go-signal (not indicated in the figure), one of the two subjects starts to exert a phasic effort to push the sphere into the assigned container, as revealed by a clear peak, slightly anticipating the load cell signal. While in panel A the blue and the red peaks clearly alternate, in panel B the trend is less clear, showing some degree of superimposition of the two traces during some of the trials. Note that in both conditions the instructions were exactly the same: "Place the sphere into the target container". The only difference between the two conditions concerned the monetary reward. In Condition 1, each member of the pair was winning at any trial, while in Condition 3, each member of the pair was winning only when the target container was the one at her right side, requiring the pulling of the sphere towards the container requiring an index finger adduction (FDI muscle not involved).

This qualitative difference between conditions is quantitatively shown in Figure 3, depicting the average values of FDI muscle EMG, recorded from each subject while pushing the sphere into the target container placed at her left side in the three experimental conditions. EMG data, after normalization, were averaged subject by subject (n=24) by pooling the last 12 trials before the signal recorded from the load cell.

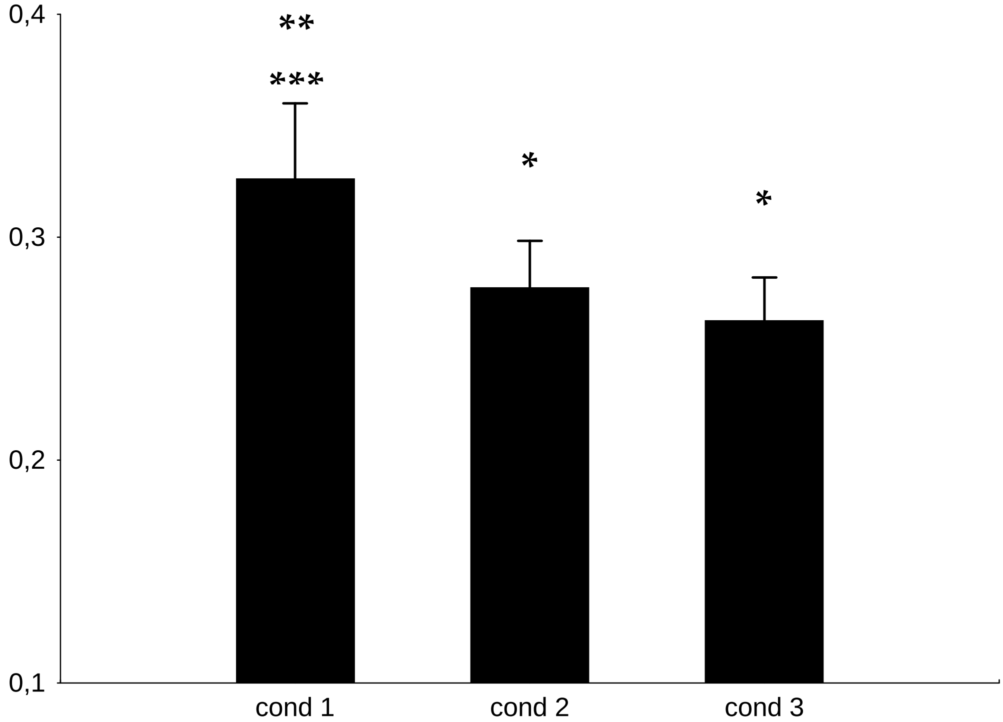

Figure 4. Mean values of EMG signals recorded from the FDI muscle for all subjects in the three experimental conditions, when pushing the sphere into the target container placed at her left side. Whiskers above each histogram depict the standard error of mean. Ordinates: *z*-score of EMG signals. Asterisks indicate the presence of a significant difference between conditions (*, difference from Condition 1; **,*** difference from Conditions 2 and 3, respectively).

An Analysis of Variance (ANOVA) was performed on the data with Experimental Condition as three levels within-subjects factor. Results showed that the factor Experimental Condition was statistically significant ($F(2,46)=4.48$, $p<0.017$). Post-hoc analysis (Newman-Keuls) revealed that Condition 1 was significantly ($p<0.05$) different from Conditions 2 and 3. This result means that the muscle activity is stronger in Condition 1 than in Conditions 2 and 3. However, as indicated in Table 1, the game outcome does not reflect this difference, and subjects, interviewed at the end of the experiment, never reported the voluntary use of different strategies in the different conditions.

Questionnaire

One of the aims of the present work was to verify if the degree of prosocial propensity modulates muscle involvement of the pushing subjects, in response to different monetary incentives among conditions. Using the questionnaire's answers, we built up three indicators (*SC1*, *SC2*, *SC3*) to sort subjects according to their attitude to coordinate and cooperate for mutual benefit (see Appendix for details). For each of these indicators subjects have been divided into two subgroups with respect to the index-related median score, defining the high- (H=above median) and the lowprosocial (L=below median) group of subjects.

Regression Behavioural Model.

To process the information gathered through the questionnaire and to control for the robustness of the results obtained with the ANOVA, we developed a regression behavioral model. Our data set is distributed along four relevant dimensions: time, trials, subjects and conditions. The potential information of this stock of data is not fully exploited by standard analysis of variance, since ANOVA does not control for many potential sources of variability, such as the muscle effort exerted by subject's couplemate, or individual fixed effects. Therefore, we considered the following dynamic multiple regression model

$$EMG_{it} = \alpha_0 + \alpha_1 EMG_{it-n} + \alpha_2 EMG_{jt-n} + \beta_2 C_2 + \beta_3 C_3 + \gamma_{it} X_{it} + \eta_i + \varepsilon_{it}, \qquad (1)$$

The dependent variable EMG_{it} is subject i's EMG signal at time t, when involved in pushing the sphere towards the target container. The righthand side of the equation models the set of explanatory variables. Specifically, EMG_{it-n} is the lagged EMG of subject i and EMG_{jt-n} is the lagged EMG of subject j (couplemate of subject i). To perform successfully the task it is required a continuous exchange of information between subjects, by the pressure exerted by their index fingers. The EMG_{it-n} variables reflect the intention of subject i to push the sphere into the target container. At the same time, since the task requires the collaboration of subject j, the lagged EMG_{jt-n} take into account that subject i's effort depends on the opposition force exerted by subject's j finger. Thus, the dynamic part of the regression model represents the motor communication between subjects i and j. Other factors that might have influenced the motor behavior of subjects could have been determined by strain or stress and learning-by-doing. To account for these factors, we introduced in vector X_{it} the time length of trials and the sequence order of trials over the entire experiment. The reason of our choice is that lengthy trials may have been more expensive in terms of attention, thus affecting the effort spent by subjects. Furthermore, subjects' effort might have been differently modulated over the course of the experiment, due to a better knowledge of her couplemate and/or to the improvement in their motor ability. Several non

observable characters of subjects (such that religion, education, family conditions etc..) may influence the dependent variable. The term η_i represents a vector of individual dummies, that control the regression model for this individuals' heterogeneity. Finally, C_2 and C_3 are two dummies for condition 2 and 3 respectively, controlling for experimental conditions instructions.
The time horizon of the regression considered 12 observations before the maximum EMG level, included. Lags in regressors EMG_{it-n} and EMG_{jt-n} have been set at 2 and 5 time periods (n=2, 5). This accounts for a period of time ranging from 80 ms (2 * 40 ms, being the sampling frequency 25 Hz) to 200 ms (5 * 40 ms). This choice was based on the observation that when a perturbation is applied during a precision grip a latency of 60-80 ms is required to increase the grip force to restore an adequate safety margin, preventing frictional slips (Eliasson et al., 1995). Thus, we defined this time range in order to include the minimal reaction time to a change in the load force applied by subject *j*, plus a possible delay determined by the fact that the grasping requires a coordination between two subjects and not only between two fingers of the same hand.

Regression Results.

The relevant estimation results are presented in table (2) below. The first column (POOL) reports the estimation results for the entire set of subjects (24). The other six columns provide results relative to each high/low prosocial sub-groups according to indicators SC1, SC2 and SC3. In particular, HSC_z and LSC_z (z=1, 2, 3) refer to High and Low prosocial individuals, respectively.

	POOL	HSC_1	LSC_1	HSC_2	LSC_2	HSC_3	LSC_3
EMG_{it-5}	0.1121	0.1439	0.0838	0.1840	0.0567	0.1560	0.0924
	(4.59)***	(4.22)***	(2.32)**	(5.78)***	(1.56)	(4.72)***	(2.46)**
EMG_{it-2}	0.2898	0.2918	0.2634	0.2634	0.2724	0.2294	0.3233
	(12.69)***	(8.49)***	(8.78)***	(7.67)***	(8.76)***	(6.73)***	(10.54)***
EMG_{jt-5}	0.0898	0.0405	0.1421	0.0473	0.1552	0.0615	0.1196
	(3.82)***	(1.38)	(3.79)***	(1.64)	(4.06)***	(2.11)**	(3.07)***
EMG_{jt-2}	0.0330	0.1153	-0.0451	0.0748	-0.0073	0.0474	0.0281
	(1.41)	(3.61)***	(1.38)	(2.46)**	(0.21)	(1.61)	(0.75)
C_2	-0.0041	-0.0095	0.0003	-0.0098	0.0040	-0.0079	0.0015

	(0.40)	(0.70)	(0.02)	(0.72)	(0.27)	(0.59)	(0.10)
C_3	-0.0244	0.0156	-0.0630	0.0105	-0.0535	-0.0093	-0.0389
	(2.44)**	(1.10)	(4.49)***	(0.73)	(3.82)***	(0.67)	(2.72)***
Constant	0.1044	0.0867	0.1322	0.1319	0.1092	0.1748	0.0632
	(4.46)***	(2.66)***	(0.3.89)***	(3.71)***	(2.80)***	(4.56)***	(2.05)**
Observations	1755	855	870	885	870	880	875
R-squared	0.3955	0.4774	0.3455	0.4276	0.4121	0.4472	0.3665

Robust t statistics in parentheses

* significant at 10%; ** significant at 5%; *** significant at 1%

Table 2: Ordinary Least Square Regressions keeping 12 observations before the maximum EMG level, included. Normalization over the entire data set.

Overall, coefficients of lagged variables are positive and significant, suggesting that each couple of subjects successfully tried to coordinate their index fingers as a pair of agonists. However, looking at the magnitude of coefficients for different groups of subjects substantial differences emerge between high-prosocial and low-prosocial individuals. In particular, the following results appear

R1) $\alpha_{1,-2}^{H,L} > \alpha_{1,-5}^{H,L}$

R2) $\alpha_{1,-2}^{H} \approx \alpha_{1,-2}^{L}$, $\alpha_{1,-5}^{H} > \alpha_{1,-5}^{L}$

R3) $\alpha_{2,-2}^{H} > \alpha_{2,-5}^{H}$, $\alpha_{2,-2}^{L} < \alpha_{2,-5}^{L}$ ($\alpha_{2,-2}^{H} > \alpha_{2,-2}^{L}$, $\alpha_{2,-5}^{H} < \alpha_{2,-5}^{L}$)

where $\alpha_{1,-n}^{H,L}$ and $\alpha_{2,-n}^{H,L}$ (*n*=2, 5) refer to coefficients of regressors EMG_{it-n} and EMG_{jt-n} respectively, while *H* and *L* apexes stay for high-prosocial and low-prosocial subjects.

For both *H* and *L* the autoregressive component of the regression model (the lagged EMG_{it-n} variables) shows that the current effort EMG_{it} of subject *i* is positively linked to her own past efforts, and that the magnitude of the coefficients decreases the farther-off are the lags (result R1). This is consistent with figure (3), which shows that intensity of muscles effort progressively increases, and reaches its peak at the instant at which the sphere is dropped. However, result R2 reveals that the EMG recording of *H* subjects display a smoother time profile than that of *L*'s.

Result R3 describes how the current reaction of subject i depends on past motor behavior of subject *j*. Overall, estimated coefficients are significantly non negative. However, looking at the size of coefficients it emerges a striking difference between *H* and *L* individuals. Current muscle effort of high-prosocial subjects is influenced mainly by the more recent behavior of their couplemates, while muscle effort of low-prosocial subjects depends only on their own behavior. Considering high-prosocial subjects, the estimated coefficients on EMG_{jt-5} are not significantly different from zero in two of the three regressions (HSC_1 and HSC_2) and significant at the 5% level but close to zero in the HSC_3 case. On the contrary, coefficients on EMG_{jt-2} are positive and significant in HSC_1 and HSC_2 and not significant in HSC_3. Exactly the reverse pattern occurs with low-prosocial subjects: coefficients on EMG_{jt-5} are significant at a 1% level, while those on EMG_{jt-2} are not significant in all cases (LSC_1, LSC_2 and LSC_3).

To suggest an interpretation for this result, one may consider as a benchmark case of perfect coordination two fingers of a single hand grasping an object to the purpose of dropping it somewhere. In this case the applied grip force is synchronically balanced to optimize the motor behavior, and therefore the pressure exerted by finger *i* is instantaneously matched with the pressure of finger *j.* In statistical terms, perfect synchronicity would be revealed by a lack of significant correlation between current effort of finger *i* and past efforts of finger *j.* In light of these considerations, the estimated coefficients on EMG_{jt-n}'s shows that on average high-prosocial subjects have been able to coordinate more efficiently with their couplemates than low-prosocial subjects.

Finally, the estimated coefficients on dummies C_2 and C_3 confirm the main results obtained with the ANOVA procedure, indicating that on average subjects exerted a lower pushing effort in condition 3 than in condition 1. Only the estimated coefficient of C_3 is negative (-0.0244) and 5% significant (*t*=2.44). However, once we distinguish between high-prosocial and low-prosocial subjects, the estimated coefficients of C_3 is negative and significant at a 1% level in the low-prosocial subsample, only. This pattern arises whatever index of social capital is used.

DISCUSSION

A conspicuous body of experimental economics literature has shown that the allocation of legitimate property rights significantly affect the strategic behavior of individuals (Hoffman and Spitzer, 1985; Hoffman et al, 1994, 1996; Ruffle, 1998; Cherry, 2001; Cherry et al, 2002; Oxoby and Spraggon, 2008). Indeed, individuals often interact, by simply committing themselves to a set of shared social norms, basically concerned with a broad view of property rights, that include not only specific entitlements on things but also on actions[3]. In this respect, property rights provide an efficient device to prompt cooperation among individuals, avoiding costly mind-reading activity. In the light of these considerations, we set up an experimental framework, aimed at investigating the effects of formal property rights, not supported by the legitimation of a specific task, when interaction is removed from any complex perspective-taking activity. Specifically, we performed an experiment with pairs of subjects, prevented from any visual or verbal exchange, engaged in a pure motor coordination game divided in three conditions, perfectly identical both from the point of view of the motor task and from that of the monetary stake. Each couple of subjects was asked to hold a small sphere between their right index fingers and to alternately drop it into one of two containers placed below their hands, while electromyography of participants' right FDI muscle was recorded. This muscle has the function to abduct the index finger, that is to draw the index finger away from the middle finger. Thus, it is the muscle more involved in pushing the sphere towards the leftmost container, while it remains relaxed when the participant is asked to place the sphere into the rightmost container by exerting a finger adduction. Our aim was to compare FDI muscle activity when participants were asked to push the sphere into the leftmost container under different rewarding schemes. In Condition 1 the completion of each trial entailed an equal prize assigned to both subjects, in Condition 2 only the subject who had pushed the sphere towards the leftmost container obtained the prize. Therefore, FDI muscle involvement in pushing the sphere was coupled with a monetary reward. Thus, in conditions 1 and 2 the rules of the game formally entitled pushing subjects to get a reward at every trial. In condition 3 the reward was given entirely to the pulling subject. Thus, FDI muscle involvement in pushing the sphere was not coupled with any monetary reward.

Subjects were able to coordinate almost perfectly across conditions 1, 2 and 3 (see Table 1). Thus, from a distributional point of view it does not emerge any difference in behavior associated with the different incentive protocols. However, substantial differences arose from EMG data processing, revealing not only that muscle involvement in executing the same motor act is affected by the allocation of formal property rights, but also that the modulation of the effort is correlated with the degree of prosocial propensity of subjects. To measure the social attitude of participants we used the answers to the questionnaire taken from Putnam's *Social Capital Benchmark Survey* to

[3] In this latter sense a property right defines a specific social role.

construct three indexes of *social capital*, that we used to split the sample of subjects into high-prosocial and low-prosocial individuals. With respect to these two groups of individuals our main result was that high-prosocial subjects performed the task without any significant difference among conditions, while low-prosocial subjects exerted a significant lower effort in Condition 3 than in Condition 1.

When a small object is gripped between the tips of the index finger and thumb and held stationary in space, the applied grip force is synchronically balanced to optimize the motor behavior. In addition, the control of the grip force is automatically influenced by the weight of the object (load force) and by a safety margin factor related to the individual subject (Westling et al. 1984, Edin et al., 1995). This is fundamental to avoid the accidental drop of the object. If we consider that the two index fingers of a pair of subjects act on the sphere as a pair of agonists, we can assume that the major effort exerted by low-prosocial individuals reflects a higher level of the safety margin factor. Our results suggest that the reallocation of property right from the 'pushing' to the 'pulling' subject modulated the safety margin factor in low-prosocial individuals, only. A likely interpretation of this effect is that low-prosocials differently evaluated the successful outcome of the pushing action in response to the reallocation of property right to the pulling subject. On the contrary, since the safety margin set by high prosocials did not change across conditions, we argue that their motor behavior did not react to changes in the distribution of property rights.

We believe that these results find place within the debate concerning the *endowment effect*. The endowment effect describes the tendency of individuals to value a good they possess more highly that the same good they do not possess. In other words the mere ownership of something causes to increase the subjective value attributed to it. A huge amount of evidence supports the relevance of this effect in actual behavior (Tversky and Kahnemann; 1981, Kahneman *et al*., 1990, 1991; Thaler; 1992, Plott and Zeile, 2003) and, recently, this phenomenon has been reported even in animals such as chimpanzees (Brosnan et al. 2007). The presence of the endowment effect has questioned the traditional assumption of rationality at the basis of behavioral models in economics and law. More specifically, if the influence of a subjective sense of ownership induces people to evaluate goods and rights irrationally, then the standard prediction of the Coase Theorem fails. This theorem claims that if transaction costs are sufficiently low, private bargaining will lead to an efficient outcome regardless of the initial allocation of property rights. The importance of this theorem is not related to market activity only, but it applies to any conflict may arise in social interaction.

Although the endowment effect is considered one of the most robust phenomenon in the emerging field of behavioral economics, it is recognized to be quit changeable, appearing and disappearing with different degrees of intensity depending on the context (Brown and Gregory; 1999, Sayman and Öncüler; 2005). Thus, as pointed out by Jones and Brosnan (2008), to investigate deeply the nature of this phenomenon, a primary goal in the research agenda should be the identification of those factors that may help to predict its appearance. We think that the evidence reported in our paper provides insight in this direction. In particular, our results suggest that the prosocial attitude might be one factor influencing the emerging of the endowment effect.

The behavior of high-prosocial individuals, revealing that the allocation of property rights doesn't modulate their muscle activity, seems to agree with the rationality principle assumed by the Coase Theorem. On the other side, the result showing that low-prosocial individuals exert different muscle effort according to the ownership of the reward, can be considered a further evidence of the endowment effect.

ACKNOWLEDGEMENTS

This work has been realized with the contribution of E.C. Grants Robot-Cub and Poeticon. We thank the *Fondazione Cassa di Risparmio* di Ferrara for their financial support.

REFERENCES

Bobo, L., delli Carpini, M., Dawson, M., Guterbock, T., Putnam, R., Rahn, W., & Thompson J. (2001). Social capital benchmark survey, *Saguaro Seminar at John F. Kennedy School of Government*, Harvard University, mimeo.

Brosnan, S., Jones, O., Lambeth, S., Mareno, M., Richardson, A., & Schapiro, S. (2007), Endowment effects in chimpanzees, *Current Biology*, 17, 1704-1707.

Brown, T., & Gregory, R. (1999). Survey: Why the WTA-WTP disparity matters, *Ecological Economics*, 28, 323-324.

Cherry, T. (2001). Mental accounting and other-regarding behaviour: evidence from the lab, *Journal of Economic psychology*, 22, 605-615.

Cherry, T., Frylblom, P., & Shogren, J. (2002). Hardnose the dictator. *American Economic Review*, 92, 1218-1221.

Edin, B., Essick, G., Trulsson, M., & Olsson, K.(1995). Receptor encoding of moving tactile stimuli in humans. I. Temporal pattern of discharge of individual low-threshold mechanoreceptors. *Journal of Neuroscience,* 15, 830-847.

Eliasson, A., Forssberg, H., Ikuta, K., Apel, I., Westling, G., & Johansson, R. (1995). Development of human precision grip. V. anticipatory and triggered grip actions during sudden loading. *Experimental brain research,* 106, 425-433.

Fahr, R., & Irlendbusch, R. (2000). Fairness as a constraint on trust in reciprocity: earned property rights in a reciprocal exchange experiment. *Economic Letters*, 66, 275-282.

Hoffman, E., & Spitzer, M. (1985), Entitlements, rights and fairness: an experimental examination of subjects concept of distributive justice. *Journal of Legal Studies*, 14, 259-297.

Hoffman, E., McCabe, K., & Smith, V. (1994). Preferences, property rights, and anonimity in bargaining games. *Games and Economic Behaviour*, 7, 346-380.

Hoffman, E. McCabe, K., & Smith, V. (1996). Social distance and other-regarding behaviour in dictator games. *American Economic Review*, 86, 653-660.

Johansson, R., & Westling G. (1984). Roles of glabrous skin receptors and sensorimotor memory in automatic control of precision grip when lifting rougher or more slippery objects. Experimental Brain Research, 3, 550-564

Jones, S. & Brosnan, F. (2008), Law, biology, and property: A new theory of the endowment effect. *William & Mary Law Review*, 49, 1935-1990.

Kahneman, D., Knetsch, J., & Thaler, R. (1990). Experimental test of the endowment effect and the Coase Theorem. *Journal of Political Economy*, 6, 1325-1348.

Kahneman, D., Knetsch, J., & Thaler, R. (1991). Anomalies: The endowment effect, loss aversion, and the status quo bias. *Journal of Economic Perspectives*, 5, 193-206.

Oxoby, R., & Spraggon, J. (2008). Mine and yours: property rights in dictator games. *Journal of Economic Behaviour and Organization*, 65, 703-713

Plott, C., & Zeile K. (2003), The willingness to pay/willingness to accept gap, the "endowment effect" and experimental procedures for eliciting valuation. *Caltech Social Science Working Paper #1732*.

Ruffle, B. (1998). More is better, but fair is fair: tipping in dictator and ultimatum games. *Games and Economic Behaviour*, 23, 247-265.

Sayman, S., & Öncüler, A. (2005). Effects of study design characteristics on the WTA-WTP disparity: A meta analytical framework. *Journal of Economics and Psychology*, 26, 290-292.

Thaler R. (1992), *The Winner's Curse*, Princeton University Press. Princeton.

Tversky, A., & Kahneman A. (1981), Loss aversion in riskless choice: a reference-dependent model. *Quarterly Journal of Economics*, 106, 1039-1061.

APPENDIX: THE QUESTIONNAIRE AND THE SOCIAL CAPITAL INDEXES

The subjects where asked to answer a questionnaire, described in detail below, designed following very closely Putnam's Social Capital Benchmark Survey (Bobo et. al., 2001). An increasing number of applications, from sociology to health economics, political science, business management, human resources and politics have used the concept of Social Capital, depending on circumstances, as synonym of rather diverse concepts, such as "generalized trust", "civic engagement", "religious belief" or "group interaction". Using the questionnaire's answers we build up indicators of these characteristics to sort our subjects according to their attitude to coordinate and cooperate for mutual benefit.

Following Bobo et al. (2001) we build up six indexes. Civic participation (cp) is constructed (see index CIVPART in Bobo et al. (2001)) as the average of three different questions, meant to measure individual involvement in civic and political activity, such as working for a political party in the past year (q1.2), attending political meetings in the past year (q1.5) and signing petitions in the past year (q1.7):

$$\boxed{cp=(q1.2+q1.5+q1.7)/3.}$$

We also build an alternative index cpext, by adding subject's answer to a specific question, namely q5.4 which asked how important was politics in their personal life (answer ranked from 4="very important" to 1="not important at all":

$$\boxed{cpext=(q1.2+q1.5+q1.7+(q5.4-1)/3)/4.}$$

Faith-based Social Capital (fbsc) is an indicator (see index FAITHBAS in Bobo et al. (2001)) constructed as the average of two questions, designed to measure participation in the life of the local religious community such as going to church in the past week (q2.8), or going to church social function in the past month (q3.6)

$$\boxed{fbsc=(q2.8+q3.6)/2.}$$

By analogy with cpext, we also consider the following:

$$\boxed{fbscext=(q2.8+q3.6+(q5.6-1)/3)/3.}$$

Organized Group Interactions (ogi) is built (see index ORGINTER in Bobo et al. (2001)) as the average of six questions, designed to measure participation in the life of the local community such as serving as an officer of some club organization in the past year (q1.1), or in a committee for some local organization in the past year (q1.2), attending a public meeting of club or civic organization in the past month (q3.7):

ogi= (q1.1+q1.3+q1.4+q3.7)/4.

Informal Group interaction (igi) is an indicator (see index SCHMOOZ in Bobo et al. (2001)) constructed as the average of six questions, designed to measure participation in the informal social network such as having friends in for the evening in the past week (q2.3); going to the home of friends in the past week (q2.4); going to club, disco, bar or place of entertainment in the past week q2.11); going to friends' house for dinner or evening in the past month (q3.4); having friends in for dinner or evening in the past month (q3.5); going to night club, disco, bar in the past month (q3.9):

isi=(q2.3+q2.4+q2.11+q3.4+q3.5+q3.9)/6

Bobo et al. (2001) also considers five additional indexes, based on social trust (STRSTCAT), group involvement without church participation (GRPINCAT), group involvement with church participation(GRP2CAT), diversity of friendship network (DIVRCAT), and composite racial group trust (RACETCAT). Due to a almost null variability in the subjects' answers (probably due to a higher homogeneity of our subject pool with respect to the relevant dimensions) we could not make any use of these additional indexes.

Finally, since we need to rank our subject pool with respect to a composite scale that would comprise all the relevant characteristics revealing attitude to coordinate and cooperate, we construct three composite measures using the indexes above outlined:

C1=(cpext+fbscext+ogi+isi)/4
C2=(cpext+fbscext+ogi)/3
C3=(cp+fbsc+ogi)/3.

THE QUESTIONNAIRE

In what follows we report the text (translated into English) of the questionnaire.

Please answer to the following questions

1. Which, if any, of these things have you done in the past year?
    - 1.1 Served as an officer of some club or organization
    - 1.2 Worked for a political party
    - 1.3 Served on a committee for some local organization
    - 1.4 Attended a public meeting on town or school affairs
    - 1.5 Attended a political rally or speech
    - 1.6 Made a speech
    - 1.7 Signed a petition
    - 1.8 Wrote a letter to the paper
    - 1.9 Wrote an article for a magazine or newspaper

2. Which, if any, of these things have you done in the past week?

    - 2.1. Discussed politics
    - 2.2. Had dinner in a restaurant
    - 2.3 Had friends in for the evening
    - 2.4 Went to the home of friends
    - 2.5 Saw a movie
    - 2.6 Made a personal long distance call
    - 2.7 Read a book
    - 2.8 Went to church
    - 2.9 Watched a sports event on TV
    - 2.10 Went out to watch a sports event
    - 2.11 Went to club, disco, bar or place of entertainment
    - 2.12 Spent time on a hobby
    - 2.13 Wrote a personal letter or e-mail
    - 2.14 Received a personal letter or e-mail

3. How many times, if any, did you do any of these activities in the past month?
    - 3.1 Made a contribution to charity
    - 3.2 Did volunteer work
    - 3.3 Donated blood
    - 3.4 Went to friends' house for dinner or evening
    - 3.5 Had friends in for dinner or evening

- 3.6 Went to church social function
- 3.7 Went to meeting of club or civic organization
- 3.8 Went to dinner at restaurant
- 3.9 Went to night club, disco, bar
- 3.10 Went to live theater, opera, concerts
- 3.11 Went to sporting event
- 3.12 Went to the movies

4. Which of the following things are part of "the good life" in your opinion?
   - 4.1 A home you own
   - 4.2 A yard and lawn
   - 4.3 A second car
   - 4.4 A vacation home
   - 4.5 A swimming pool
   - 4.6 A happy marriage
   - 4.7 No children
   - 4.8 One or two children
   - 4.9 A job that pays more than average
   - 4.10 A job that is interesting
   - 4.11 A job that contributes to the welfare of society
   - 4.12 College education for my children
   - 4.13 Travel abroad
   - 4.14 A second color TV set
   - 4.15 Really nice clothes
   - 4.16 A lot of money

5) For each of the following, indicate how important it is in your life. Would you say it is:

1. Very important
2. Rather important
3. Not very important
4. Not at all important
5. I don't know

- 5.1 A home you own Family
- 5.2 A yard and lawn Friends
- 5.3 A second car Leisure time
- 5.4 Politics
- 5.5 Work
- 5.6 Religion
- 5.7 Service to others

6) Taking all things together, would you say you are:

4. Very happy

3. Quite happy
2. Not very happy
1. Not at all happy
0. Don't know

7) With which of these two statements do you tend to agree? (CODE ONE ANSWER ONLY)
A. Regardless of what the qualities and faults of one's parents are, one must always love and respect them
B. One does not have the duty to respect and love parents who have not earned it by their behavior and attitudes
7.1 Tend to agree with statement A
7.2 Tend to agree with statement B
7.3 Don't know

8) Generally speaking, would you say that most people can be trusted or that you need to be very careful in dealing with people?
8.1 Most people can be trusted
8.2 Need to be very careful
8.3 Don't know

9) Do you think most people would try to take advantage of you if they got a chance, or would they try to be fair?
9.1 Would take advantage
9.2 Would try to be fair
9.3 Don’t know

Current Biology 19, 381–385, March 10, 2009 ©2009 Elsevier Ltd All rights reserved DOI 10.1016/j.cub.2009.01.017

Report

The Motor Somatotopy of Speech Perception

Alessandro D'Ausilio,[1] Friedemann Pulvermüller,[2] Paola Salmas,[3] Ilaria Bufalari,[1] Chiara Begliomini,[1] and Luciano Fadiga[1,3,*]
[1]DSBTA
Section of Human Physiology
University of Ferrara
Ferrara 44100
Italy
[2]Cognition and Brain Sciences Unit
Medical Research Council
Cambridge CB2 7EF
UK
[3]IIT, The Italian Institute of Technology
Genova 16163
Italy

Summary

Listening to speech recruits a network of fronto-temporo-parietal cortical areas [1]. Classical models consider anterior (motor) sites to be involved in speech production whereas posterior sites are considered to be involved in comprehension [2]. This functional segregation is challenged by action-perception theories suggesting that brain circuits for speech articulation and speech perception are functionally dependent [3, 4]. Although recent data show that speech listening elicits motor activities analogous to production [5–9], it's still debated whether motor circuits play a causal contribution to the perception of speech [10]. Here we administered transcranial magnetic stimulation (TMS) to motor cortex controlling lips and tongue during the discrimination of lip- and tongue-articulated phonemes. We found a neurofunctional double dissociation in speech sound discrimination, supporting the idea that motor structures provide a specific functional contribution to the perception of speech sounds. Moreover, our findings show a fine-grained motor somatotopy for speech comprehension. We discuss our results in light of a modified "motor theory of speech perception" according to which speech comprehension is grounded in motor circuits not exclusively involved in speech production [8].

Results

Recent years have seen a major change in views about the function of motor and premotor cortex [11]. Once believed to be an output system, slavishly following the dictate of the perceptual brain, the motor brain is now recognized as critical component of perceptual and cognitive functions. This challenges the classical sensory versus motor separation [12]. Similarly, traditional models of language brain organization separated perceptual and production modules in distinct areas [1, 2]. However, a large amount of data is accumulating against the reality of such a strict anatomo-functional segregation [5–9, 13, 14]. The motor theory of speech perception (MTSP) [3], an early precursor of a new zeitgeist, most radically postulated that the articulatory gestures, rather than sounds, are critical for both production and perception of speech (see [4]). On neurobiological grounds, fronto-temporal circuits are thought to play a functional role in production as well as comprehension of speech. The coactivation of motor circuits and the concurrent perception of self-produced speech sounds during articulations might lead to correlated neuronal activity in motor and auditory systems, triggering long-term plastic processes based on Hebbian learning principles [15–17]. The postulate of a critical role of actions in the formation of speech circuits is paralleled in more general action-perception theories emphasizing a critical role of action representations in action-related perceptual processes [18]. However, a majority of researchers are still skeptical toward a general role of motor systems in speech perception, admitting, if at all, only a subsidiary role of motor areas and reserving the critical role to superior temporal and inferior parietal cortices [19].

A recent series of studies directly investigated the activities in motor areas during speech perception. Passive listening to phonemes and syllables was shown to activate motor [5–8] and premotor [9] areas. Interestingly, these activations were somatotopically organized according to the effector recruited in the production of these phonemes [5, 6, 8] and in accordance with motor activities in overt production [8, 9]. A distinctive feature of action-perception theories in general and in the domain of language specifically is that motor areas contribute to perception [4, 16, 20]. However, all the above mentioned studies are inherently correlational, and it has been argued that in absence of a stringent determination of a causal role played by motor areas in speech perception, no final conclusion can be drawn in support of motor theories of speech perception [10]. The only empirical evidence in favor of this view is represented by a recent repetitive TMS study suggesting that ventral premotor cortex (PMv) may play some role in phonological discrimination [21]. In our view, however, this study fails to offer a convincing proof of the causal influence that motor areas may exert. Because of the spread and the variety of possible effects elicited by a 15 min TMS stimulation, such an offline rTMS protocol might have indeed modified the activity of a larger network of areas, possibly including posterior receptive language centers [22]. Moreover, there is no evidence of an effector-specific effect, i.e., that stimulating tongue representation induced specific deficits in the perception of tongue-related phonemes.

Here, we set out to investigate the functional contributions of the motor-articulatory systems to specific speech-perception processes. To this end, a cross-over design orthogonalizing the effect of brain-phonology concordance with those of linguistic stimuli and TMS loci was chosen. Phonemes produced with different articulators (lip-related: [b] and [p]; tongue-related: [d] and [t]) were presented in a phoneme-discrimination task. The effect of TMS to lip and tongue representations in precentral cortex, as previously described by fMRI [8], was investigated. Double TMS pulses were applied just prior to stimuli presentation to selectively prime the cortical activity specifically in the lip (LipM1) or tongue

*Correspondence: fdl@unife.it

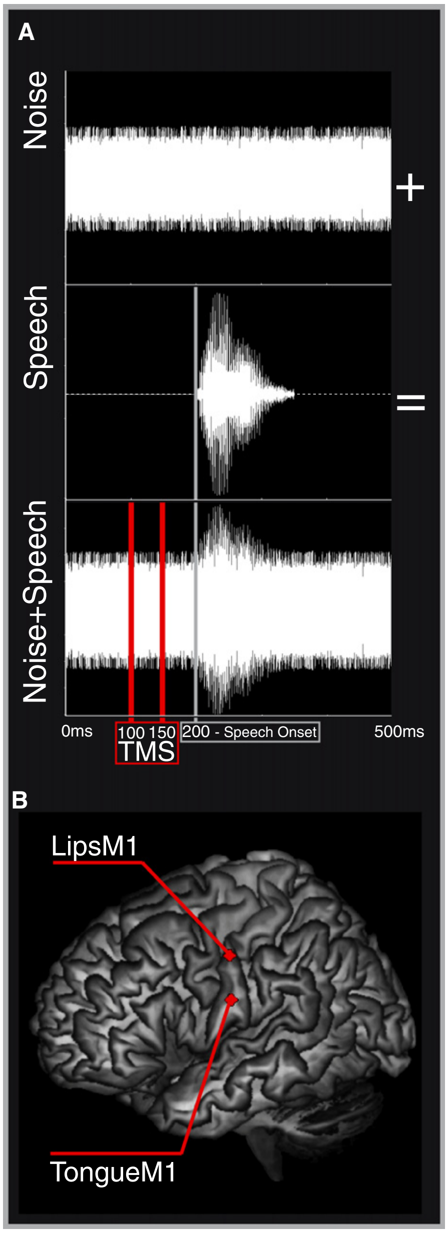

Figure 1. Stimuli, TMS Timing, and Regions of Stimulation

(A) Noise, speech sound, and experimental stimulus waveforms. Noise and speech recordings were mixed into a single trace. TMS (vertical red lines) was applied in double pulses 100 and 150 ms after noise onset. Speech sounds started 200 ms after noise onset (gray vertical line).

(B) LipM1 and TongueM1 normalized mean coordinates are projected on a standard template [8, 34].

(TongueM1) area (Figure 1). We hypothesized that focal stimulation would facilitate the perception of the concordant phonemes ([d] and [t] with TMS to TongueM1), but that there would be inhibition of perception of the discordant items ([b] and [p] in this case). Behavioral effects were measured via reaction times (RTs) and error rates.

RT performance showed a behavioral double dissociation between stimulation site and stimulus categories (Figure 2). RT change of phonological decisions induced by TMS pulses to either the TongueM1 or LipM1 showed opposite effects for tongue- and lip-produced sounds. The interaction of the phoneme type and stimulation site factors was significant ($F[1,36] = 17.578$; $p < 0.0005$), and the post-hoc analysis evidenced a significant difference between labial ([b], [p]) and dental ([d], [t]) phonemes for each of the stimulation sites. As hypothesized, recognition of lip-produced phonemes was indeed faster than that of tongue-produced ones when stimulating the LipM1 (labial = 94.8% ± 5.3% SEM; dental = 117.3% ± 3.7% SEM; $p = 0.009$), and the stimulation of the TongueM1 induced the reverse pattern (labial = 113.6% ± 6.4% SEM; dental = 93% ± 5.1% SEM; $p = 0.024$). In addition, labial and dental stimuli recognition was faster when stimulating their concordant M1 representation compared with that to the discordant stimulation locus (labial, $p = 0.015$; dental, $p = 0.009$). Therefore, the stimulation of a given M1 representation led to better performance in recognizing speech sounds produced with the concordant effector compared with discordant sounds produced with a different effector. These results provide strong support for a specific functional role of motor cortex in the perception of speech sounds.

In parallel, we tested whether TMS was able to modulate the direction of errors (Figure 3). Errors were grouped in two classes: lip-phoneme errors (L-Ph-miss) and tongue-phoneme errors (T-Ph-miss). The ANOVA showed a significant interaction effect ($F[1,36] = 4.426$; $p < 0.05$). Post-hoc comparisons

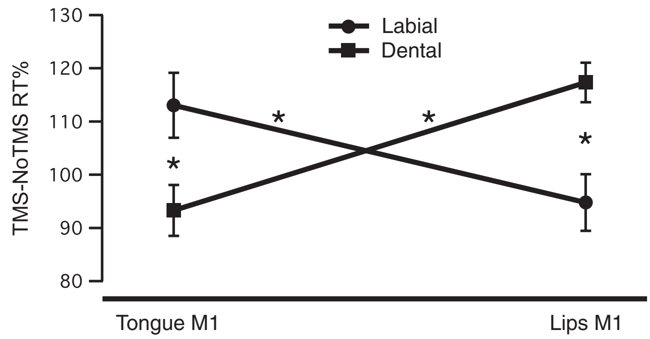

Figure 2. Reaction Times during Speech Discrimination

Effect of TMS on RTs show a double dissociation between stimulation site (TongueM1 and LipM1) and discrimination performance between class of stimuli (dental and labial). The y axis represents the amount of RT change induced by the TMS stimulation. Bars depict SEM. Asterisks indicate significance ($p < 0.05$) at the post-hoc (Newman-Keuls) comparison.

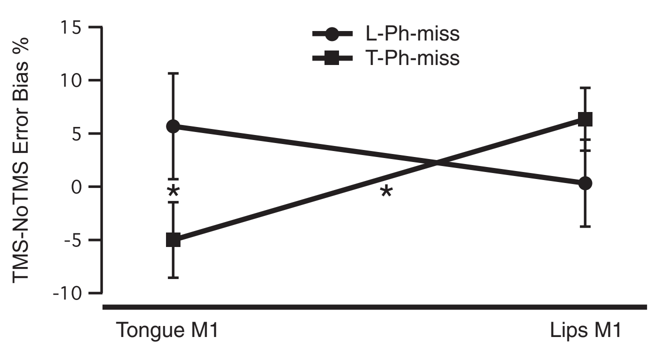

Figure 3. Accuracy Results

We tested whether TMS was able to modulate the direction of errors, i.e., if the stimulation of the TongueM1 increases the number of labial sounds erroneously classified as dental and vice versa. After TMS, a dissociation between stimulation site (TongueM1 and LipM1) and kind of errors (L-Ph-miss, T-Ph-miss) was found. The y axis represents the amount of error change induced by the TMS stimulation. Other conventions as in Figure 2.

revealed more L-Ph-miss than T-Ph-miss errors when stimulating the TongueM1 representation (p = 0.049), and also more T-Ph-miss errors when stimulating the LipM1 relative to the TongueM1 (p = 0.012). Therefore, the error pattern confirmed the dissociation already seen in the RT data. As a matter of fact, the stimulation of a given motor representation led to a perceptual bias in favor of speech sounds concordant with the stimulation site. Stimulation of the tongue area made lip sounds tend to be perceived as dentals and, vice versa, lip area TMS made [d] and [t] sound like bilabials.

Discussion

The double dissociation we found in the present work provides evidence that motor cortex contributes specifically to speech perception. As shown by both RTs and errors, the perception of a given speech sound was facilitated by magnetically stimulating the motor representation controlling the articulator producing that sound, just before the auditory presentation. Inhibitory effects were seen for discordant speech sounds. Computationally speaking, our stimulation might preactivate, or prime, a given M1 sector by increasing the excitability of neurons therein. This higher excitability might lead to faster RTs if that area contributes to a task. The reduction of performance observed in the other class of stimuli can be explained by mechanisms of lateral inhibition between competing representations. Similarly, TMS-induced priming of one specific representation may bias the system toward activating the already preactivated representation, leading to the observed error pattern. The direction of our effects suggests that our TMS protocol is enhancing the activity in M1 locally, in agreement with other results reported in TMS literature [23–25] and with the work by Pulvermüller and colleagues [26] describing a similar effect at the semantic level. Two factors might have caused facilitation of subjects' behavioral performance: (1) TMS timing and (2) basal cortical activity. In fact, a single TMS pulse disrupts cortical processing for a limited time window, by synchronizing neuronal activities. Animal models actually hold that inhibition turns into facilitation after a short time window depending on pulse strength [27]. Alternatively, the direction of effects can be accounted by cortical state dependency. It's well known that motor thresholds vary according to the cortico-spinal basal activity (i.e., muscle contraction). Analogously, the recent work of Silvanto and colleagues [28, 29] showed that TMS induces both behavioral facilitation and inhibition according to the basal activity of the target cortical area. Although both explanations are equally probable, the double dissociation we found is the strongest proof to support our central hypothesis. It should be stressed, however, that our finding does not prove that M1 is directly involved in speech perception. A possible explanation for the facilitation of the perception of phonemes motorically congruent with the stimulated site is that the synchronous excitation of M1 neurons induced by TMS may have exerted in turn the facilitation of neurons located in premotor areas, somatotopically connected with M1 through bidirectional cortico-cortical links.

Biologically grounded models of speech and language have previously postulated a functional link between motor and perceptual representations of speech sounds [8, 30]. We demonstrate here for the first time a specific causal link for features of speech sounds. The relevant areas in motor cortex seem to be also relevant for controlling the tongue and lips, respectively. As mentioned, the MTSP [3, 4] had postulated a critical role of articulatory phonetic gestures for the perception of speech. However, this theory also claimed a modular status of the linguistic phoneme system, which was thought to be functionally dissociated from the nonlinguistic motor system. This position is difficult to reconcile with the finding of congruency between cortical areas for speech perception, articulation, and nonlinguistic movements of tongue or lips [8]. Here we therefore provide only partial support for the MTSP, and we propose that the motor gestures critical for speech perception are processed by the same brain parts and circuits involved also in the production of other, nonlinguistic, movements. We are not suggesting, however, that the motor cortex is an area for phonological discrimination per se; rather, we favor the idea that it might be part of a larger network. This latter claim is also supported by a large number of studies showing an integrated brain network for speech processing as opposed to a single localized module [1, 13, 14, 16, 19, 31]. We propose that TMS of M1 might have unbalanced the network dynamics of action-perception circuits, likely involving motor, premotor, and temporo-parietal areas.

The present results might be of potential interest in the rehabilitation of aphasia. Current experimental protocols, showing initial exciting results, are indeed evaluating the possible benefit of repeated TMS (rTMS) application in these patients [31]. TMS is typically used to trigger (or inhibit) plastic processes in conjunction with standard rehabilitation protocols. However, rTMS effects spread uncontrollably to other areas, eventually resulting in a global functional reshaping of whole-brain dynamics. Event-related TMS, such as the one presented in our study, might be potentially more spatially selective and thus more effective. Single pulses or short trains might in fact be more efficient in triggering local plastic processes in selected neuronal populations. We therefore propose that innovative rehabilitation programs based on recent neuroscientific findings about action-perception circuits [13, 16], such as the intensive language-action therapy [32], in conjunction with event-related TMS protocols might be more effective also at chronic stages of aphasia.

Experimental Procedures

Subjects

Ten healthy right-handed subjects volunteered after giving informed consent and were paid for their participation (mean age, 26.07; SD, 2.91; 6 female). None had any history of neurological disease, trauma, or psychiatric syndrome and all had normal hearing. Procedures were approved by the local ethical committee.

Stimuli

Subjects listened to one out of four stimuli in each trial: [b], [p], [d], and [t] spoken before a [œ] sound, through headphones. [b] and [p] are labial sounds, requiring the critical lip movement for their production, whereas [d] and [t] are dental sounds that require a significant tongue movement. Each stimulus was the vocal recording of an actor. In order to avoid ceiling effects in the phoneme identification task, we immersed vocal recordings in 500 ms of white noise. Each vocal stimulus was presented 200 ms after the beginning of white noise. The noise/stimulus ratio was set in a pilot experiment (11 subjects) to let subjects respond correctly $\approx$75% of cases. Task RT and accuracy, grouped into labial (mean RT, 839 $\pm$ 59.95 ms SEM; mean accuracy, 77.58% $\pm$ 5.36% SEM) and dental (mean RT, 815 $\pm$ 53.11 ms SEM; mean accuracy, 75.76% $\pm$ 5.78% SEM) sounds did not differ significantly in the pilot experiments (RT, $t(10) = 1.249$; $p = 0.24$; accuracy, $t(10) = 0.234$; $p = 0.82$).

Task

Subjects were asked to listen and recognize the consonants and respond as fast as possible with a four-button pad. Buttons were configured in a diamond shape and the relative position, with associated consonant letters, was shown during the experiment on a screen in front of them. Responses were given with the left index finger. Response recording, stimuli presentation, and TMS triggering were controlled by a custom-made Basic script running under the MS-DOS environment to warrant timing accuracy.

TMS

TMS stimulation was delivered through a figure-of-eight 40 mm coil and a Magstim Rapid stimulator (Magstim, Whitland, UK). The 25 mm coil was used to allow a more focal stimulation. First *Dorsal Interosseus* (FDI) mapping and resting motor threshold (rMT) evaluation was assessed by using standard protocols [33]. Motor-evoked potentials (MEP) were recorded by using a standard tendon-belly montage with Ag/Cl electrodes. Signal was band-pass filtered (50–1000 Hz) and digitized (5 kHz). TongueM1 and LipM1 localization were, instead, based on standardized coordinates with respect to the functionally defined best stimulation site of the FDI muscle. Specifically, for lip and tongue area stimulation, we chose the mean MNI coordinates corresponding to the peak motor cortex activation probability (t-values), during lip and tongue movements and articulation, revealed by a previous fMRI study (lips: −56, −8, 46; tongue: −60, −10, 25; Figure 1A; [8]). In parallel, also FDI MNI coordinates were taken from the literature (−37, −25, 58; [34]). In the following step, MNI coordinates (FDI, TongueM1, and LipM1) were first transformed into the 10–20 EEG system space (Münster T2T-Converter: http://wwwneuro03.uni-muenster.de/ger/t2tconv/conv3d.html) and then the distance between FDI/tongue and FDI/lip were calculated in the same standard space. Therefore, TongueM1 and LipM1 were located according to differential 10–20 EEG coordinates centered on the functionally defined FDI location. In each subject, the FDI was first functionally mapped, and then TongueM1 and LipM1 were located according to the differential 10–20 EEG coordinates (lips: 6.6% of *nasion-inion* distance in the anterior direction and 5.8% of the inter-tragus distance in the lateral direction–FDI mean distance: 5.5 cm; tongue: 8.6% anterior and 11.6% lateral–FDI mean distance: 3.35 cm; mean distance between lips and tongue: 2.15 cm). In the stimulated trials, two pulses with 50 ms interval were delivered at 110% of the FDI rMT. Coil orientation was maintained at 45° with respect to the interhemispheric fissure. Pulses were given 100 ms and 150 ms after noise onset; thus, the last TMS pulse occurred 50 ms prior to consonant presentation (see Figure 1).

Procedure

Subjects first completed a block of trials with no TMS intervention, to train participants and to test their ability to accomplish the task up to our criteria ($\approx$75% of correct trials; a total of 60 trials, 15 each stimulus category). Upon successful completion of this learning phase, they were entered in the TMS mapping block. The right FDI primary motor representation was located and marked on the left hemisphere, and the rMT was measured. LipM1 and TongueM1 representations were marked on the scalp with respect to the functionally defined FDI spot (for the procedure see the TMS section). After the mapping session, two blocks were presented in succession separated by a 2 min interval. TMS stimulation over LipM1 and TongueM1 was delivered in different blocks, whose order was counterbalanced across subjects. In each block, subjects had to complete 80 trials, 60 with TMS and 20 random catch trials. Random catch trials were exactly the same as the TMS trials except that no TMS was applied. Catch trials were used as a reference to evaluate the effect induced by TMS on behavior.

Measures and Analysis

Experimental measures included RTs and errors. RTs were calculated from the beginning of consonant sound presentation (200 ms after noise onset). The RT data were collapsed into two categories: labial and dental sounds. Preliminary analyses showed that there were no significant differences between the voiced ([b], [d]) and unvoiced ([p], [t]) phonemes. Subjects' performance was normalized by computing the percentage of variation of mean RT in TMS-stimulated trials with respect to trials without TMS. Errors were considered as the amount of responses erroneously attributed to the other category (misses). Errors were collapsed in two categories according to their falling in the other group of stimuli (L-Ph-miss and T-Ph-miss). Single subjects' error scores were expressed as the percentage of change between stimulated trials and TMS-free ones. Separate analyses of variance (ANOVA) were conducted on RT and error data, including the factors phoneme type (labial versus dental or, in the error analysis, L-Ph-miss versus T-Ph-miss) and stimulation site (LipM1 versus TongueM1). Significant interactions were further investigated with Newman-Keuls post-hoc comparisons (alpha = 0.05).

Acknowledgments

L.F. is supported by Italian Ministry of Education; by the E.C. grants Contact, Robot-cub, and Poeticon; and by Fondazione Cassa di Risparmio di Ferrara. F.P. is supported, in part, by MRC grants (U1055.04.003.00001.01, U1055.04.003.00003.01) and by the E.C. grant Nestcom. L.F., F.P., and A.D. conceived the experiment and cowrote the paper. A.D., P.S., and I.B. acquired the data. A.D., C.B., P.S., and I.B. analyzed the data.

Received: November 16, 2008
Revised: January 3, 2009
Accepted: January 5, 2009
Published online: February 12, 2009

References

1. Gernsbacher, M.A., and Kaschak, M.P. (2003). Neuroimaging studies of language production and comprehension. Annu. Rev. Psychol. *54*, 91–114.
2. Damasio, A.R., and Geschwind, N. (1984). The neural basis of language. Annu. Rev. Neurosci. *7*, 127–147.
3. Liberman, A.M., Cooper, F.S., Shankweiler, D.P., and Studdert-Kennedy, M. (1967). Perception of the speech code. Psychol. Rev. *74*, 431–461.
4. Galantucci, B., Fowler, C.A., and Turvey, M.T. (2006). The motor theory of speech perception reviewed. Psychon. Bull. Rev. *13*, 361–377.
5. Fadiga, L., Craighero, L., Buccino, G., and Rizzolatti, G. (2002). Speech listening specifically modulates the excitability of tongue muscles: a TMS study. Eur. J. Neurosci. *15*, 399–402.
6. Watkins, K.E., Strafella, A.P., and Paus, T. (2003). Seeing and hearing speech excites the motor system involved in speech production. Neuropsychologia *41*, 989–994.
7. Pulvermüller, F., Shtyrov, Y., and Ilmoniemi, R.J. (2003). Spatiotemporal patterns of neural language processing: an MEG study using minimum-norm current estimates. Neuroimage *20*, 1020–1025.
8. Pulvermüller, F., Huss, M., Kherif, F., Moscoso del Prado Martin, F., Hauk, O., and Shtyrov, Y. (2006). Motor cortex maps articulatory features of speech sounds. Proc. Natl. Acad. Sci. USA *103*, 7865–7870.
9. Wilson, S.M., Saygin, A.P., Sereno, M.I., and Iacoboni, M. (2004). Listening to speech activates motor areas involved in speech production. Nat. Neurosci. *7*, 701–702.
10. Toni, I., de Lange, F.P., Noordzij, M.L., and Hagoort, P. (2008). Language beyond action. J. Physiol. (Paris) *102*, 71–79.
11. Rizzolatti, G., and Luppino, G. (2001). The cortical motor system. Neuron *31*, 889–901.
12. Young, R.M. (1970). Mind, Brain and Adaptation in the Nineteenth Century. Cerebral Localization and Its Biological Context from Gall to Ferrier (Oxford: Clarendon Press).
13. Pulvermüller, F. (2005). Brain mechanisms linking language and action. Nat. Rev. Neurosci. *6*, 576–582.
14. Skipper, J.I., Nusbaum, H.C., and Small, S.L. (2006). Lending a helping hand to hearing: another motor theory of speech perception. In Action to Language via the Mirror Neuron System, M.A. Arbib, ed. (Cambridge: Cambridge University Press), pp. 250–285.
15. Fry, D.B. (1966). The development of the phonological system in the normal and deaf child. In The Genesis of Language, F. Smith and G.A. Miller, eds. (Cambridge, MA: MIT Press), pp. 187–206.
16. Pulvermüller, F. (1999). Words in the brain's language. Behav. Brain Sci. *22*, 253–336.
17. Braitenberg, V., and Pulvermüller, F. (1992). Entwurf einer neurologischen Theorie der Sprache. Naturwissenschaften *79*, 103–117.
18. Rizzolatti, G., and Craighero, L. (2004). The mirror-neuron system. Annu. Rev. Neurosci. *27*, 169–192.
19. Hickok, G., and Poeppel, D. (2007). The cortical organization of speech processing. Nat. Rev. Neurosci. *8*, 393–402.
20. Pulvermüller, F., and Preissl, H. (1991). A cell assembly model of language. Network-Comp. Neural. *2*, 455–468.
21. Meister, I.G., Wilson, S.M., Deblieck, C., Wu, A.D., and Iacoboni, M. (2007). The essential role of premotor cortex in speech perception. Curr. Biol. *17*, 1692–1696.
22. Matsumoto, R., Nair, D.R., LaPresto, E., Najm, I., Bingaman, W., Shibasaki, H., and Lüders, H.O. (2004). Functional connectivity in the human language system: a cortico-cortical evoked potential study. Brain *127*, 2316–2330.

23. Töpper, R., Mottaghy, F.M., Brügmann, M., Noth, J., and Huber, W. (1998). Facilitation of picture naming by focal transcranial magnetic stimulation of Wernicke's area. Exp. Brain Res. *121*, 371–378.
24. Grosbras, M.H., and Paus, T. (2003). Transcranial magnetic stimulation of the human frontal eye field facilitates visual awareness. Eur. J. Neurosci. *18*, 3121–3126.
25. Hayward, G., Goodwin, G.M., and Harmer, C.J. (2004). The role of the anterior cingulate cortex in the counting Stroop task. Exp. Brain Res. *154*, 355–358.
26. Pulvermüller, F., Hauk, O., Nikulin, V.V., and Ilmoniemi, R.J. (2005). Functional links between motor and language systems. Eur. J. Neurosci. *21*, 793–797.
27. Moliadze, V., Zhao, Y., Eysel, U., and Funke, K. (2003). Effect of transcranial magnetic stimulation on single-unit activity in the cat primary visual cortex. J. Physiol. *553*, 665–679.
28. Silvanto, J., Muggleton, N., and Walsh, V. (2008). State-dependency in brain stimulation studies of perception and cognition. Trends Cogn. Sci. *12*, 447–454.
29. Silvanto, J., Cattaneo, Z., Battelli, L., and Pascual-Leone, A. (2008). Baseline cortical excitability determines whether TMS disrupts or facilitates behavior. J. Neurophysiol. *99*, 2725–2730.
30. Wennekers, T., Garagnani, M., and Pulvermüller, F. (2006). Language models based on Hebbian cell assemblies. J. Physiol. (Paris) *100*, 16–30.
31. Devlin, J.T., and Watkins, K.E. (2007). Stimulating language: insights from TMS. Brain *130*, 610–622.
32. Pulvermüller, F., and Berthier, M.L. (2008). Aphasia therapy on a neuroscience basis. Aphasiology *22*, 563–599.
33. Rossini, P.M., Barker, A.T., Berardelli, A., Caramia, M.D., Caruso, G., Cracco, R.Q., Dimitrijević, M.R., Hallett, M., Katayama, Y., Lücking, C.H., et al. (1994). Non-invasive electrical and magnetic stimulation of the brain, spinal cord and roots: basic principles and procedures for routine clinical application. Report of an IFCN committee. Electroencephalogr. Clin. Neurophysiol. *91*, 79–92.
34. Niyazov, D.M., Butler, A.J., Kadah, Y.M., Epstein, C.M., and Hu, X.P. (2005). Functional magnetic resonance imaging and transcranial magnetic stimulation: effects of motor imagery, movement and coil orientation. Clin. Neurophysiol. *116*, 1601–1610.

Blackwell Publishing EJN European Journal of Neuroscience

Force requirements of observed object lifting are encoded by the observer's motor system: A TMS-study

Journal:	*European Journal of Neuroscience*
Manuscript ID:	draft
Manuscript Type:	Research Report
Date Submitted by the Author:	
Complete List of Authors:	Alaerts, Kaat; Katholieke Universtiteit Leuven, Biomedical Kinesiology Senot, Patrice; University of Ferrara, Dep. S.B.T.A., Sect. of Human Physiology Swinnen, Stephan; Katholieke Universtiteit Leuven, Biomedical Kinesiology Craighero, Laila; University of Ferrara, Dep. S.B.T.A., Sect. of Human Physiology Wenderoth, Nicole; Katholieke Universtiteit Leuven, Biomedical Kinesiology Fadiga, Luciano; University of Ferrara, Dep. S.B.T.A., Sect. of Human Physiology
Key Words:	Transcranial magnetic stimulation, mirror neurons, movement, motor system, perception

Proposed Journal section: Cognitive Neuroscience

Associate Editor: Douglas P. Munoz

Title: Force requirements of observed object lifting are encoded by the observer's motor system: A TMS-study

Running title: Encoding of observed grip force in motor cortex

Authors:

Alaerts Kaat[1a] & Senot Patrice[2a], Swinnen Stephan P[1], Craighero Laila[2b], Wenderoth Nicole[1b], Fadiga Luciano[2b]

a) These authors contributed equally in designing, performing, and analysing the experiments as well as writing the manuscript.

b) These authors contributed equally in developing the research question, designing and supervising the experiments and writing the manuscript.

Authors addresses:

[1] Research Center of Movement Control and Neuroplasticity

Department of Biomedical Kinesiology

Group Biomedical Sciences

Katholieke Universtiteit Leuven

Tervuursevest 101 - B-3001 Heverlee, BELGIUM

[2] D.S.B.T.A. Section of Human Physiology

Faculty of Medicine

University of Ferrara

Via Fossato di Mortara 17/19 - 44100 Ferrara, ITALY

Corresponding authors:

[1] **Wenderoth Nicole**

Motor Control Laboratory

Research Center of Movement Control and Neuroplasticity

Department of Biomedical Kinesiology,

Group Biomedical Sciences, Katholieke Universtiteit Leuven, Belgium

Tervuursevest 101 B-3001 Heverlee, BELGIUM

e-mail: Nici.Wenderoth@faber.kuleuven.be

Phone: +32 16 32 91 57

Fax: +32 16 32 91 97

[2] **Fadiga Luciano**

D.S.B.T.A. Section of Human Physiology

Faculty of Medicine, University of Ferrara, Italy

Via Fossato di Mortara 17/19 - 44100 Ferrara, ITALY

e-mail: fdl@unife.it

Phone: +39-0532-455338

Fax: +39-0532-455242

Number of pages: 23

Number of figures: 4

Number of tables: 0

Number of words

- **(i) whole manuscript: 6570**
- **(ii) the Abstract: 248**
- **(iii) the Introduction: 589**

Keywords: Mirror system, Action observation, Transcranial magnetic stimulation, Grip force

ABSTRACT

Several transcranial magnetic stimulation (TMS) studies report that viewing other's actions facilitates the neural representation site of the onlooker's muscles that are recruited during the actual execution of the observed action. With the present study, we investigated whether this muscle-specific facilitation of the observer's motor system reflects the degree of muscular force that is exerted in an observed action.

Two separate TMS-experiments are reported in which corticomotor excitability was measured in the hand area of the primary motor cortex (M1) while subjects observed the lifting of objects with different weights. The type of action 'grasping and lifting the object' was always identical but the grip force varied according to the object's weight.

In accordance to previous findings, activity of M1 was shown to modulate in a muscle-specific way, such that only those parts of M1 that control the specific muscles used in the observed lifting action, become increasingly facilitated. Moreover, the muscle-specific facilitation pattern of M1 was shown to modulate in accordance to the force requirements of the observed actions, such that corticomotor excitability was considerably higher for observing heavy object lifting compared to light object lifting. Overall, these results indicate that observed object grasping, requiring different force levels, is mirrored onto the observer's motor system in a highly muscle-specific manner, as measured in M1.

The measured force-dependent modulations of corticomotor activity in M1 are hypothesised to be functionally relevant for the observer's ability to infer the observed grip force and consequently the weight of the lifted object.

INTRODUCTION

In social interactions, humans demonstrate the remarkable ability to understand and interpret the behaviour of other people. Recently, neuroscience is increasingly focussing on the role of the observer's motor system during action understanding (Rizzolatti & Craighero 2004). This possibility comes from results of single cell recordings in monkeys demonstrating the existence of "mirror neurons", which were shown to respond both when a monkey performs a certain action and when it observes another person performing the same action (Di Pellegrino et al., 1992). In humans, several neuroimaging and neurophysiological studies have identified the inferior frontal gyrus (IFG) as well as the parietal cortex to be key areas of the 'human mirror neuron system' (Grafton et al., 1996;Decety et al., 1997;Cochin et al., 1998;Buccino et al., 2001;Grezes et al., 2003;Buccino et al., 2004;Lui et al., 2008). In addition, with transcranial magnetic stimulation (TMS) it was shown that parts of primary motor cortex (M1) that control particular muscles become increasingly facilitated during the mere observation of actions involving these muscles (Fadiga et al., 1995;Strafella & Paus 2000). Moreover, M1-excitability modulations reflect specific characteristics of observed actions: Next to the robust finding that modulations of M1 are strongly muscle-specific (Borroni & Baldissera 2008;Alaerts et al., 2009a;Alaerts et al., 2009b), previous research also showed that M1 activations are highly synchronized to the temporal dynamics of an observed movements (Gangitano et al., 2001;Borroni et al., 2005;Montagna et al., 2005) and lateralized to the contra-lateral hemisphere when right- versus left-hand actions are observed (Aziz-Zadeh et al., 2002). As such, it appears that visual-motor matching during observation is a highly specified process in which different features of the observed actions are encoded by the observer's motor system.

All of the above parameters (muscular involvement, temporal dynamics, used effector) can be easily derived from robust differences in the kinematics of the observed movement. However, until now, it is largely unclear whether features which are less salient in the kinematic signal, such as the force requirements of an observed lifting action, are also matched to the observer's motor system. Some behavioural studies already indicated that the weight of a box (and consequently the force needed to lift it) can be inferred quite accurately by observing another person lifting it (Runeson & Frykholm 1981;Bingham 1987), and interestingly, some recent experiments demonstrated that the observer's motor system might be involved in this task. More specifically, it was shown that the active lifting of weights interferes with concomitant weight judgement tasks (Hamilton et al., 2004). These findings suggest that similar force-related parameters are coded both in the motor plan and in the action representation evoked by the observation of the action. A following study localized the sites of interaction between perceptual and motor processes in several frontal and parietal areas and particularly in the IFG and M1 (Hamilton et al., 2006). However, activations in both regions were not confirmed during the mere observation of weight lifting (Hamilton & Grafton, 2007). As such, their actual involvement in weight perception needs to be established further.

As TMS is known to be very powerful in assessing activity modulations at the level of M1, the present study used this technique to explore possible force-related activity modulations in the observer's motor system during the observation of lifting objects with different weights.

In two separate experiments, performed in two laboratories, cortical excitability was measured in M1 during the observation of lifting objects with different weights. Thus, the type of action 'grasping and lifting the object' was always identical but the grip force varied according to the object's weight.

EXPERIMENT 1

Observation of object lifting with a precision grip

Experiment 1 was designed to test whether force requirements are encoded in the observer's motor system during observation of an actor lifting two different objects of explicit different weights, using a precision grip. Experiment 1 was run in Ferrara, Italy.

MATERIALS & METHODS

Subjects. Eight subjects (5 males, 3 females) with age between 20 and 32 (mean: 22) participated after providing informed consent. All experimental protocols were approved by the local ethics board in accordance to The Code of Ethics of the World Medical Association (Declaration of Helsinki) (Rickham 1964). All participants were right-handed, as assessed with the Edinburgh Handedness Questionnaire (Oldfield 1971) and were naive about the purpose of the experiment.

Electromyographic recordings and TMS. Surface electromyography (EMG) was performed with Ag-AgCl electrodes placed according to a belly-tendon configuration. EMG activity was recorded from the right First Dorsal Interosseous (FDI) finger muscle, an intrinsic hand muscle acting as agonist for precision grip.

Focal transcranial magnetic stimulation (TMS) was performed by means of a 70 mm figure-of-eight coil connected to a Magstim 200 stimulator (Magstim). The coil was positioned over the left hemisphere, tangentially to the scalp with the handle pointing backward and medially at 45° away from the mid-sagittal line, such that the induced current flow was in a posterior-anterior direction, i.e. approximately perpendicular to the central sulcus. An articulated arm (Manfrotto, Italy) was used to keep coil position during the experiment. The optimal scalp position was defined as the position from which Motor Evoked Potentials (MEPs) with maximal amplitude were recorded in the right FDI muscle. The rest motor threshold (rMT) was defined as the lowest stimulus intensity evoking MEPs in the right FDI with an amplitude of at least 50µV in 5 out of 10 consecutive stimuli (Rossini et al., 1994). Stimulation intensity was set at 120% of the rest motor threshold for all experimental trials. EMG signals were band-pass filtered (50-1000Hz), digitized (2000Hz) and stored on a computer for off-line analysis.

General Procedure. Participants were seated comfortably on a dentist like armchair, their arms stretched out on a arm rest and their hand lying relaxed and pronated. They faced a small stage with black floor and background. A square metallic platform aligned with the subject's sagittal plane supported the target object on which action was performed. The actor was seated fully visible on the front right of the participant and acted with his right hand on the target object, parallel to the subject's frontal plane. The actor reached to grasp the object with his right hand, lifted it, held it few seconds and then replaced it at its initial position. The two objects presented (Figure 1A) were of different shape and explicitly of different weight despite they both could be grasped by opposing the tips of the thumb and index finger (precision grip) thanks to a common handle. The first object ("Light") was a 10g piece of ribbon cable that was held erected by individualizing the wires at the lower extremity of the ribbon. The other object ("Heavy") was a 500g brass balance weight with a handle made of the same ribbon cable used for the Light object. In each trial, the actor's hand initially lied pronated on the table, pushing with the fingertips a hidden switch placed at about 20 cm from the object (Figure 1B). One of the two objects was then placed on the platform. A vocal warning ("pronto") was provided to signal the incoming of a new trial. The contact time of the actor's fingers with the object, and the lifting latency were provided respectively by an electric circuit switched on by the contact between both fingers and the object's handle and switched off by the separation between the object and the metallic platform. Each of the two objects was presented 15 times with presentation order randomized within subjects. A single TMS pulse was delivered at random time during the lifting phase of the observed movement, approximately 1.3 sec after movement onset (Figure 1B). In total, 30 MEPs were recorded for each subject. Before the experimental session, subjects could see the objects and were allowed to experience their respective weight.

Data reduction and analysis. From the EMG data, peak-to-peak amplitudes of the MEPs were determined. Since EMG background activation is known to modulate MEP amplitude (Hess et al., 1987;Devanne et al., 1997), pre-stimulation EMG was assessed in both experiments by computing root-mean-square error scores (RMSe) across a 50 ms interval prior to the TMS stimulation. For each subject and for each muscle separately, mean and standard deviation of the EMG background scores were computed over all trials. Trials for which EMG background was above the mean + 2.5 standard deviation were removed from the analysis. Trials for which the MEP amplitude was inferior to the mean EMG background were also discarded. Finally, extreme peak-to-peak amplitudes values in the remaining trials were removed from the analysis under the following criteria: outliers were considered as values larger than Q3+1.5 x (Q3-Q1) with Q1 the first quartile and Q3 the third quartile computed over the whole set of trials for each subject. Following these three criteria one subject was discarded, due to 80% of bad trials in one of the observation conditions. From the remaining subjects, 13 % of trials were discarded in total.

For each subject, MEP amplitudes recorded for each observation condition were then normalized relative to the subjects' maximal MEP amplitude (measured over all trials and conditions). Subsequently, normalized MEPs were averaged among subjects. RMSe scores of each condition were also normalized relative to the maximal RMSe score (measured over all trials and conditions) and averaged among subjects.

Statistics. Paired T-tests were used to compare peak-to-peak MEP amplitude data recorded during the observation of the heavy and light weight lifting. Similar statistical analyses were applied to the background EMG data (normalized RMSE-scores) to assess whether the MEP amplitude scores were confounded by modulations in background EMG.

RESULTS

During the **observation** of object lifting with a precision grip, individual normalized MEP amplitudes in the FDI muscle revealed a systematic modulation relative to the weight of the lifted object. For six out of seven subjects, MEP amplitude scores were higher during observation of heavy object lifting compared to light object lifting, and this difference was significant in one subject [S3, t=2.395, p=.038] (Figure 2A). At the group level (n=7), this consistent trend led to significantly higher normalized MEP amplitudes for the heavy compared to the light weight observation condition [t=2.8, p=.031] (Figure 2B).

A paired T-test computed on the background EMG data (normalized RMSe-scores) confirmed that the EMG background was not significantly different in the two conditions [t=.972, p=.369], indicating that experimental results are not likely explained by a modulation in background EMG.

EXPERIMENT 2

Execution and observation of object lifting with a whole hand grip

Experiment 2 was designed to test if the results found in experiment 1 are consistently found during observation of other types of grip, and if the weight-related modulation is specific for the muscles involved in action execution or if it reflects an unspecific activation of the motor system. In addition, in experiment 2, the muscle activation pattern for real execution of lifting different object weights was assessed and compared to the corticomotor responses obtained during the observation of the same lifting actions.

The main differences between the two observational paradigms regard: (i) the observation of a precision grip (Exp 1) or of a whole hand prehension (Exp 2); (ii) the recording of one muscle only (First Dorsal Interosseus) (Exp 1) or of three muscles (Opponens Pollicis, Flexor and Extensor Carpi Radialis) (Exp 2); (iii) the comparison of two weights (Exp 1) or of three weights (Exp 2), and (iv) the involvement of a real agent performing the movement (Exp 1) or the use of videos (Exp 2). Although both types of stimuli (i.e., real actions or video-taped actions) are known to induce a reactivation of primary motor cortex, it may be worth mentioning that reactivations were shown to be more salient for observing real actions compared to video-taped actions (Jarvelainen et al., 2001).

MATERIALS & METHODS

Execution

Task. Five subjects (age range 23-30; 3 females, 2 males) were instructed to observe a video displaying a grasp-lift action and to simultaneously perform the same action in synchrony with the video. The video showed the whole hand grasp and lift of drinking bottles with three different weights i.e., an empty (0 kg), a half full (1 kg) and a full (2 kg) bottle (Figure 3).

EMG. During execution, EMG was simultaneously recorded from the right Opponens Pollicis (OP) muscle and Flexor (FCR) and Extensor Carpi Radialis (ECR) muscles.

Data analysis. Each subject performed the three actions 15 times. In 12 additional trials, the EMG was recorded during maximal voluntary contraction (MVC) of each muscle. EMG-changes (amplitudes) were calculated for a short time-interval of 40 ms during the lifting of the bottle (Figure 3). EMG changes were expressed as the percentage of subjects' muscle-specific MVC-scores.

Observation

Subjects. Twelve subjects (3 males and 9 females) with an age range of 21-35 (mean: 23) participated after providing informed consent. The subjects participating in the action observation experiment were not the same subjects that participated in the action execution experiment.

Electromyographic recordings and TMS. EMG and TMS procedures were similar to those described in Experiment 1. However, MEPs were evoked and measured from the right Opponens Pollicis (OP) muscle and Flexor (FCR) and Extensor (ECR) Carpi Radialis muscles. Although stimulation settings were prioritised for the OP muscle, simultaneous measurements from the FCR and ECR are assumed to be satisfactorily similar, due to the partial overlap of representations of finger and forearm flexor and extensor muscles (Scheiber MH 1990). For all experimental trials, stimulation intensity was set at 130% of the rest motor threshold of the OP muscle. EMG signals were sampled at 5000 Hz, (CED Power 1401, Cambridge Electronic Design, UK) amplified, band-pass filtered (30-1500 Hz), and stored on a PC for off-line analysis. Signal Software (2.02 Version, Cambridge Electronic Design, UK) was used for TMS triggering and EMG recordings.

General Procedure. Participants were seated in a comfortable chair in front of a Dell P992 monitor (resolution, 1024 × 768 pixels; refresh frequency 60 Hz) on which video clips (Audio-Video Interleaved (AVI)) were displayed with a frame rate of 25 Hz (or frames per seconds). The experimental video clips showed the target object and the model's right hand that acted upon it. The model's hand entered the scene from the subject's right side, reached to grasp the object and subsequently lifted it out of the scene in the vertical plane (Figure 3). The three target objects were plastic drinking bottles with a weight of respectively 0 kg (empty), 1 kg (half full) and 2 kg (full). All bottles were grasped with a whole hand grip i.e. by using the thumb and hand palm (Figure 3). Additionally, a control video clip was presented to the subjects showing only an empty white background without any overt action (i.e. Baseline). All video clips lasted for 10 seconds. Each of the 4 video clips was presented 20 times in blocks of four, with the block presentation order randomized within and across subjects. During the presentation of each video clip, a single TMS pulse was delivered at a random time point during the bottle lifting phase (Figure 3). Video presentation timing was controlled by Blaxton Video Capture software (South Yorkshire, UK). In total, 80 MEPs were recorded from each subject. Before the experimental session, all video clips were presented to the subjects in order to familiarize them with the experimental stimuli. During the session, they were instructed to keep their hands and forearms as relaxed as possible and to pay full attention to the video presented, such that they could report the type of video after each trial.

Data reduction and analysis. The same procedures as in Experiment 1 were adopted for data analysis. Following this procedure, only 4% of all trials (of all subjects) were discarded from further analyses for each muscle (OP-FCR-ECR).

Statistics. MEP Amplitude data recorded during the observation of the three experimental video clips, were subjected to a two-way analysis of variance (ANOVA) with repeated measures, with the within factors 'Muscle' (OP, FCR, ECR) and 'Grip force' (Empty, Half Full, Full). All significant interactions were analysed further using Fisher LSD post-hoc tests (Statistica 7.0, StatSoft. Tulsa, USA).

Similar statistical analyses were applied to the background EMG data (normalized RMSE-scores) to assess whether the MEP amplitude scores were confounded by modulations in background EMG.

RESULTS

During the **execution** of object lifting with a whole hand grip, OP and ECR muscles were found to be more involved in the action compared to the FCR muscle (normalized EMG muscle activity recorded from the OP and FCR are visualized in Figure 4A). This was revealed by the two-way ANOVA interaction between 'Muscle' (OP, FCR, ECR) and 'Grip Force' (Empty, Half Full, Full) [$F(4,16)=3.22$, $p<.05$]. Main effects of 'Muscle' [$F(2,8)=4.91$, $p<.05$] and 'Grip force' [$F(2,8)=55.68$, $p<.001$] were also found. Post-hoc analysis of the two-way interaction revealed that for the OP and ECR, all force levels were significantly different from one another, and that modulations in grip force - related to the weight of the lifted object - were more pronounced in the OP and ECR muscle, compared to the FCR muscle (see Figure 4A).

During the **observation** of object lifting, normalized MEP amplitudes were shown to modulate systematically with the force requirements of the action. Moreover, force-related modulations in MEP responses were exclusively found for muscles involved in the execution of the observed action. This was revealed by the two-way ANOVA interaction between 'Muscle' (OP, FCR, ECR) and 'Grip Force' (Empty, Half Full, Full) [$F(4,44)=3.46$, $p<.05$]. A main effect of 'Muscle' [$F(2,22)=3.81$, $p<.05$], but not of 'Grip force' [$F(2,22)=2.37$, $p=.117$] was also found. Post-hoc analysis of the two-way interaction revealed that MEP responses evoked from the OP muscle were significantly higher for observing the lifting of the half full or full bottle compared to observing the empty bottle [both, $p<.01$] (Figure 4B). MEP scores yielded from the ECR muscle showed a similar modulation (Empty: 0.48 ± 0.02; Half full: 0.51 ± 0.03; Full: 0.53 ± 0.02) [Empty versus Half full, $p=.05$; Empty versus Full, $p=.007$]. In the FCR, on the other hand, no differences in MEP scores were measured for observing the different weight lifting [$p>.2$] (Figure 4B).

The background EMG was generally small and condition-specific modulations were minimal. This was tested by conducting a similar two-way ANOVA analysis (within factors 'Muscle' and 'Grip Force') to the corresponding background EMG data (normalized RMSE-scores). None of the main or interaction effects reached significance [all $F<1.5$, $p>.21$], which indicated that the MEP amplitude scores were not confounded by modulations in background EMG.

GENERAL DISCUSSION

With the present TMS-experiments, we tested whether the observer's motor system reflects the force requirements of observed actions. Our results indicated that, in accordance to previous findings, corticomotor modulation during action observation is specific for those muscles involved in the execution of the observed action, and that this muscle specific modulation is influenced by the force requirements of the observed actions, such that higher corticomotor excitability was found for the heavy object conditions than for the light object conditions.

Perception of object lifting activates the human motor system in a force-related way

In two separate experiments, carried out in two distinct laboratories, we examined whether the force requirements of an observed action are encoded in the observer's motor system during the process of visual-motor matching. Addressing the same research question, the two experiments differed mutually according to some set-up related aspects. First, in experiment 1, *live* actions were presented to the observing subjects, whereas in experiment 2, *video* presentation was used. Second, although both experiments presented (right-hand) 'grasp-lift' actions of different object weights, Exp 1 showed a '*precision grip*' (i.e., opposing the tips of the thumb and index finger), whereas Exp 2 showed a '*whole hand grip*' (i.e., using the thumb and hand palm). Consequently, the type of the 'to be grasped objects' also differed, particularly with respect to the weight ranging from 0 to 500g in Exp 1, and from 0 over 1000 to 2000g in Exp 2. However, despite these differences, both experiments established the same robust results, namely a facilitation of the observer's motor system which corresponded to the force requirements of the observed lifting actions. Experiment 2 additionally confirmed that the force-related facilitation of M1 was highly specific to the actual muscles used in the observed lifting actions. In this view, we extend previous findings on the properties of this system by showing that the level of grip force is represented in the observer's motor system. Thus, observation-to-execution mapping includes also some dynamical features of motor control, such as grip force.

The actual **execution** of successful grasps and lifts of objects involves several neuronal mechanisms, some of them being concerned with fine-tuning the grip force of the grasping fingers, and others with the transformation of object properties into motor actions (Castiello 2005). In this respect, the IFG is suggested to be involved in selecting the most appropriate 'motor prototype', such as the type of grip that is effective in interacting with the target object (Fagg & Arbib 1998), whereas the actual fine-tuning of grip force has been shown to rely strongly on primary motor cortex activity (Muir & Lemon 1983;Lang & Schieber 2004). Interestingly, there are several indications that similar brain areas may be involved during the mere perception of object lifting. Indeed, a number of studies convincingly demonstrated that the IFG is not only involved during action execution, but also during the mere observation of actions, such that it is considered to be a key area of the human mirror neuron system (Rizzolatti et al., 1996;Grafton et al., 1996;Nishitani & Hari 2000;Johnson-Frey et al., 2003;Fazio et al., 2009). More specifically, in the context of observing the lifting of different weights, an elegant study by Pobric (2006) demonstrated that perceptual weight judgements depends significantly on activity within the IFG, i.e., disruptive rTMS at this site impaired judgements of the weight of a box lifted by another person, but not judgements on the weight of a bouncing ball, and rTMS at a control site did not have this effect (Pobric & Hamilton 2006). Consistently, a study using functional magnetic resonance imaging (fMRI) also identified the IFG as well as M1 to be involved during perceptual weight judgement (Hamilton et al., 2006). Considering that IFG is strongly connected to M1 (Shimazu et al., 2004;Dum & Strick 2005), it can be argued that the measured force-dependent facilitation of M1 is a result of cortico-cortical projections from IFG mirror neurons. However to date, the actual role of M1 in the context of movement observation is still debated. On the one hand, M1 might simply be "co-activated" with IFG, thereby representing the same information as IFG. Alternatively however, it is argued that M1 plays a more functional role in movement observation by translating and representing the observed movement features in terms of muscle-related coordinates (Kilner & Frith 2007;Lepage et al., 2008;Pineda 2008;Alaerts et al., 2009b). Therefore, in relation to the studies cited above, we suggest that, in the present experiment, IFG might be occupied with representing 'motor prototypes' (such as the type of grip), whereas M1 is occupied with translating this information into 'movements', i.e., to map the types of recruited muscles as well as the level of force they produce.

Importantly, activity in IFG during weight judgement seems to rely predominantly on the general 'kinematics' of observed lifting actions, and not on object-related information about to-be-grasped objects (lifted boxes were identical in the study of Pobric et al., (2006)). However, in the present experiments, the force-related modulation in M1 might also have been triggered by 'object-related' cues, such as the filling degree of the bottle in experiment 2, or the type of material in experiment 1. Indeed, the weight of an object, and consequently the grip force needed to lift it, can quite accurately be estimated based on prior knowledge on characteristics of objects (Johansson 1998). In this context, a series of fMRI studies by Grafton & Hamilton identified the parietal region of the mirror neuron system, namely the IPS, to be a key area in representing different target objects during action observation (Hamilton & Grafton, 2006; Grafton & Hamilton, 2007). However, the observation of different object weights alone appeared to be inefficient in triggering force-related modulations in M1 (Leuven group, preliminary TMS work), suggesting that 'the action upon the object' is necessary to elicit weight/force-related responses in M1. Another ongoing study (Ferrara group) also suggests that the force-related modulation found here is more dependent on 'motion-related' cues than on explicit or implicit object-related cues. Future experiments should confirm the relative contribution of object information and purely motion-related features in mediating the force-related responses.

Nonetheless, our data convincingly showed that the motor system is recruited during observed object lifting and that its activity reflects a muscle specific force-related modulation. The potential role of M1 in representing the muscle- and force-related aspects of the observed movement has some functional significance that will be discussed in the following part.

Functional significance of force-related activity modulations in M1

Although perception and action were traditionally considered to be two distinct processes, a number of studies, using a variety of techniques, demonstrated 'mirror' activity in motor areas during the mere perception of others' action (Rizzolatti & Craighero 2004;Fadiga et al., 2005). However, to date, different hypotheses exist concerning the role of this observation-to-execution matching system.

On the one hand, it is proposed that mirror neurons contribute solely to motor planning or action preparation. Under this 'motor' hypothesis, activation of motor areas during movement observation is principally a motor resonant phenomenon (Jacob & Jeannerod 2005). However, the most widely accepted hypothesis argues that mirror neurons provide a representation of actions that allows the observer to 'automatically resonate' to observed actions in his own motor repertoire system, in order to understand or interpret the actions made by others (Gallese et al., 1996;Iacoboni et al., 2005;Craighero et al., 2007;Rizzolatti & Fabbri-Destro 2008). In accordance to the latter hypothesis, we hypothesise that the reported force-dependent modulations of M1 activity may be functionally relevant for inferring ('understanding') the grip force that is produced in the observed lifting actions (which in turn may contribute to the subjects' ability to estimate the weight of the lifted object). Unfortunately however, no formal assessment of the individual subjects' ability to judge the produced grip force was obtained such that no firm conclusions can be drawn on this point.

Aside its functional relevance in action understanding, the finding of force-related M1 activations may reveal further insights on *how accurate* force requirements are mapped within the observer's motor system. From the execution experiment (exp. 2) it was shown that the muscle activity in the OP and ECR muscle was substantially higher for lifting a full compared to lifting a half full bottle. However, the elicited corticomotor responses differed only moderately between the 'full' and 'half full' observation condition and this difference did not reach statistical significance (experiment 2). This finding suggests that force encoding was more accurately represented during movement execution than movement observation, particularly when relatively large forces were applied, Similar results were revealed by a weight discrimination study whereby subjects observed grasp/lift actions of small objects with a weights range of 50 – 850g (increasing with steps of 200g) (Hamilton et al., 2004). Even though the objects' weights were discriminated successfully, responses were fitted best by a quadratic regression, suggesting a ceiling effect for judging the highest weights (850 g). As such it can be tentatively hypothesised that a similar ceiling effect is reflected by M1 facilitation when considerably 'high' grip forces were observed. However, it should be noted that other tasks such as weight discrimination based on the observation of whole-body lifting actions did not exhibit a similar ceiling effect. Instead, a linear relationship was found when lifting actions were observed for weights ranging from 3 to 27 kg (increasing with steps of 6 kg) (Runeson & Frykholm, 1981). These differences may relate to the fact that the optimal weight discrimination range might be different between muscles which develop relatively 'weak' maximal contractions (such as distal hand muscles involved in fine-grip force tuning) and muscles developing considerably 'stronger' maximal contractions (such as proximal arm muscles involved in whole body lifting actions). Nevertheless, future studies should be conducted to specifically address this hypothesis.

In summary, the present study provides some exciting new evidence that resonant activity in motor areas is highly specified to map different features of observed actions. More specifically, data convincingly indicated that observation-induced facilitation of the observer's primary motor cortex reflects the muscular requirements of the observed movement, not only in terms of the muscle used in the observed motion, but also in terms of the force that is produced in the particular muscle.

ACKNOWLEDGEMENTS

Experiments run in Leuven, Belgium were supported by grants from the Flanders Fund for Scientific Research (Projects G.0292.05, G.0577.06 & G.0749.09) and by a grant provided the Interuniversity Attraction Poles program of the Belgian federal government (P6/29).

Experiments run in Ferrara, Italy were supported by Italian Ministry of Education; by the E.C grants Contact, Robot-cub, and Poeticon; and by Fondazione Cassa di Risparmio di Ferrara

REFERENCES

Alaerts,K., Heremans,E., Swinnen,S.P. & Wenderoth,N. (2009a) How are observed actions mapped to the observer's motor system? Influence of posture and perspective. *Neuropsychologia*, **47**, 415-422.

Alaerts,K., Swinnen,S.P. & Wenderoth,N. (2009b) Is the human primary motor cortex activated by muscular or direction-dependent features of observed movements? *Cortex*. In press

Aziz-Zadeh,L., Maeda,F., Zaidel,E., Mazziotta,J. & Iacoboni,M. (2002) Lateralization in motor facilitation during action observation: a TMS study. *Exp Brain Res.*, **144**, 127-131.

Bingham,G.P. (1987) Kinematic form and scaling: further investigations on the visual perception of lifted weight. *J Exp Psychol Hum Percept Perform.*, **13**, 155-177.

Borroni,P. & Baldissera,F. (2008) Activation of motor pathways during observation and execution of hand movements. *Soc Neurosci.*, **3**, 276-288.

Borroni,P., Montagna,M., Cerri,G. & Baldissera,F. (2005) Cyclic time course of motor excitability modulation during the observation of a cyclic hand movement. *Brain Res.*, **1065**, 115-124.

Buccino,G., Binkofski,F., Fink,G.R., Fadiga,L., Fogassi,L., Gallese,V., Seitz,R.J., Zilles,K., Rizzolatti,G. & Freund,H.J. (2001) Action observation activates premotor and parietal areas in a somatotopic manner: an fMRI study. *Eur J Neurosci.*, **13**, 400-404.

Buccino,G., Lui,F., Canessa,N., Patteri,I., Lagravinese,G., Benuzzi,F., Porro,C.A. & Rizzolatti,G. (2004) Neural circuits involved in the recognition of actions performed by nonconspecifics: An fMRI study. *J Cogn Neurosci.*, **16**, 114-126.

Castiello,U. (2005) The neuroscience of grasping. *Nat.Rev.Neurosci.*, **6**, 726-736.

Cochin,S., Barthelemy,C., Lejeune,B., Roux,S. & Martineau,J. (1998) Perception of motion and qEEG activity in human adults. *Electroencephalogr Clin Neurophysiol.*, **107**, 287-295.

Craighero,L., Metta,G., Sandini,G. & Fadiga,L. (2007) The mirror-neurons system: data and models. *Prog Brain Res.*, **164**, 39-59.

Decety,J., Grezes,J., Costes,N., Perani,D., Jeannerod,M., Procyk,E., Grassi,F. & Fazio,F. (1997) Brain activity during observation of actions - Influence of action content and subject's strategy. *Brain*, **120**, 1763-1777.

Devanne,H., Lavoie,B.A. & Capaday,C. (1997) Input-output properties and gain changes in the human corticospinal pathway. *Exp.Brain Res.*, **114**, 329-338.

Dum,R.P. & Strick,P.L. (2005) Frontal lobe inputs to the digit representations of the motor areas on the lateral surface of the hemisphere. *J Neurosci.*, **25**, 1375-1386.

Fadiga,L., Craighero,L. & Olivier,E. (2005) Human motor cortex excitability during the perception of others' action. *Curr Opin Neurobiol.*, **15**, 213-218.

Fadiga,L., Fogassi,L., Pavesi,G. & Rizzolatti,G. (1995) Motor facilitation during action observation: a magnetic stimulation study. *J Neurophysiol.*, **73**, 2608-2611.

Fagg,A.H. & Arbib,M.A. (1998) Modeling parietal-premotor interactions in primate control of grasping. *Neural Netw.*, **11**, 1277-1303.

Fazio,P., Cantagallo,A., Craighero,L., D'Ausilio,A., Roy,A., Pozzo,T., Calzolari,F., Granieri,E. & Fadiga,L. (2009) Encoding of Human Action in Broca's area. *Brain*, In press.

Gallese,V., Fadiga,L., Fogassi,L. & Rizzolatti,G. (1996) Action recognition in the premotor cortex. *Brain*, **119**, 593-609.

Gangitano,M., Mottaghy,F.M. & Pascual-Leone,A. (2001) Phase-specific modulation of cortical motor output during movement observation. *Neuroreport*, **12**, 1489-1492.

Grafton,S.T., Arbib,M.A., Fadiga,L. & Rizzolatti,G. (1996) Localization of grasp representations in humans by positron emission tomography .2. Observation compared with imagination. *Exp Brain Res.*, **112**, 103-111.

Grafton,S.T. & Hamilton,A.F. (2007) Evidence for a distributed hierarchy of action representation in the brain. *Hum Mov Sci*, **26**, 590-616.

Grezes,J., Armony,J.L., Rowe,J. & Passingham,R.E. (2003) Activations related to "mirror" and "canonical" neurones in the human brain: an fMRI study. *Neuroimage*, **18**, 928-937.

Hamilton A., & Grafton S. (2007) The motor hierarchy:from kinematics to goals and intentions. In Y Rosetti, M Kawato & P Haggard (eds), *Attention & Performance* 22

Hamilton,A., Wolpert,D. & Frith,U. (2004) Your own action influences how you perceive another person's action. *Curr.Biol.*, **14**, 493-498.

Hamilton,A.F. & Grafton,S.T. (2006) Goal representation in human anterior intraparietal sulcus. *J Neurosci.*, **26**, 1133-1137.

Hamilton,A.F., Wolpert,D.M., Frith,U. & Grafton,S.T. (2006) Where does your own action influence your perception of another person's action in the brain? *Neuroimage*, **29**, 524-535.

Hess,C.W., Mills,K.R. & Murray,N.M. (1987) Responses in small hand muscles from magnetic stimulation of the human brain. *J Physiol*, **388**, 397-419.

Iacoboni,M., Molnar-Szakacs,I., Gallese,V., Buccino,G., Mazziotta,J.C. & Rizzolatti,G. (2005) Grasping the intentions of others with one's own mirror neuron system. *PLoS Biol.*, **3**, 529-535.

Jacob,P. & Jeannerod,M. (2005) The motor theory of social cognition: a critique. *Trends Cogn Sci*, **9**, 21-25.

Jarvelainen,J., Schurmann,M., Avikainen,S. & Hari,R. (2001) Stronger reactivity of the human primary motor cortex during observation of live rather than video motor acts. *Neuroreport*, **12**, 3493-3495.

Johansson,R.S. (1998) Sensory input and control of grip. *Novartis.Found.Symp.*, **218**, 45-59.

Johnson-Frey,S.H., Maloof,F.R., Newman-Norlund,R., Farrer,C., Inati,S. & Grafton,S.T. (2003) Actions or hand-object interactions? Human inferior frontal cortex and action observation. *Neuron*, **39**, 1053-1058.

Kilner,J.M. & Frith,C.D. (2007) A possible role for primary motor cortex during action observation. *Proc Natl Acad Sci U S A.*, **104**, 8683-8684.

Lang,C.E. & Schieber,M.H. (2004) Reduced muscle selectivity during individuated finger movements in humans after damage to the motor cortex or corticospinal tract. *J Neurophysiol.*, **91**, 1722-1733.

Lepage,J.F., Lortie,M. & Champoux,F. (2008) Action-coding neurons in primary motor cortex: Making sense of M1 activity during action perception. *J Neurosci.*, **28**, 1995-1996.

Lui,F., Buccino,G., Duzzi,D., Benuzzi,F., Crisi,G., Baraldi,P., Nichelli,P., Porro,C.A. & Rizzolatti,G. (2008) Neural substrates for observing and imagining non-object-directed actions. *Soc Neurosci.*, **3**, 261-275.

Montagna,M., Cerri,G., Borroni,P. & Baldissera,F. (2005) Excitability changes in human corticospinal projections to muscles moving hand and fingers while viewing a reaching and grasping action. *Eur J Neurosci.*, **22**, 1513-1520.

Muir,R.B. & Lemon,R.N. (1983) Corticospinal neurons with a special role in precision grip. *Brain Res.*, **261**, 312-316.

Nishitani,N. & Hari,R. (2000) Temporal dynamics of cortical representation for action. *Proceedings of the National Academy of Sciences of the United States of America*, **97**, 913-918.

Oldfield,R.C. (1971) The assessment and analysis of handedness: the Edinburgh inventory. *Neuropsychologia*, **9**, 97-113.

Pineda,J.A. (2008) Sensorimotor cortex as a critical component of an 'extended' mirror neuron system: Does it solve the development, correspondence, and control problems in mirroring? *Behav.Brain Funct.*, **4**, 47.

Pobric,G. & Hamilton,A.F. (2006) Action understanding requires the left inferior frontal cortex. *Curr.Biol.*, **16**, 524-529.

Rickham,P.P. (1964) Human experimentation. Code of Ethics of the World Medical Association. Declaration of Helsinki. *BMJ.*, **2**, 177.

Rizzolatti,G. & Craighero,L. (2004) The mirror-neuron system. *Annu Rev Neurosci.*, **27**, 169-192.

Rizzolatti,G. & Fabbri-Destro,M. (2008) The mirror system and its role in social cognition. *Curr.Opin.Neurobiol.*, **18**, 179-184.

Rizzolatti,G., Fadiga,L., Gallese,V. & Fogassi,L. (1996) Premotor cortex and the recognition of motor actions. *Cogn Brain Res.*, **3**, 131-141.

Rossini,P.M., Barker,A.T., Berardelli,A., Caramia,M.D., Caruso,G., Cracco,R.Q., Dimitrijevic,M.R., Hallett,M., Katayama,Y., Lucking,C.H., Denoordhout,A.L.M., Marsden,C.D., Murray,N.M.F., Rothwell,J.C., Swash,M. & Tomberg,C. (1994) Noninvasive Electrical and Magnetic Stimulation of the Brain, Spinal-Cord and Roots - Basic Principles and Procedures for Routine Clinical-Application. *Electroencephalogr Clin Neurophysiol.*, **91**, 79-92.

Runeson,S. & Frykholm,G. (1981) Visual-Perception of Lifted Weight. *J Exp Psychol Hum Percept Perform.*, **7**, 733-740.

Scheiber MH (1990) How might the motor cortex individuate movements. *Trends Neurosci.*, **13**, 440-445.

Shimazu,H., Maier,M.A., Cerri,G., Kirkwood,P.A. & Lemon,R.N. (2004) Macaque ventral premotor cortex exerts powerful facilitation of motor cortex outputs to upper limb motoneurons. *J Neurosci.*, **24**, 1200-1211.

Strafella,A.P. & Paus,T. (2000) Modulation of cortical excitability during action observation: a transcranial magnetic stimulation study. *Neuroreport*, **11**, 2289-2292.

FIGURE LEGENDS

Figure 1.

Experimental stimuli of experiment 1.

A. Picture of the two objects grasped and lifted in front of the subject. The 500g “Heavy” object (left) was a typical brass balance weight. The 10g “Light” (right) object was a piece of ribbon cable. Both objects were grasped using the same grip hand shape.

B. Illustration of events sequence during observation of the reach-grasp-lift action executed upon the heavy object: the actor started hand pronated, then reached to the object, grasped it with precision grip, lifted it and held it over the table during 1s. A TMS pulse was delivered during the lifting phase. Time-line provides the averaged intervals (mean ± std, n=7) between the main task events (button release, hand-object contact, lift onset, TMS pulse) for action upon both the heavy and the light object.

Figure 2.

Experimental stimuli of experiment 2.

The experimental video clips showed a reach-grasp-lift action of a plastic drinking bottle with three different weights, i.e. an empty (0 kg), a half full (1 kg) and a full (2 kg) bottle. The actor entered the scene from the right side, reached to the object, grasped it with a whole hand grip and lifted it out of the scene in the vertical plane. TMS pulses were delivered at random time points during the bottle lifting phase.

Figure 3.

Results of experiment 1

A. Individual mean MEP traces for the “Heavy” (black) and “Light” (grey) experimental conditions. Asterisks indicate significant differences ($p<0.05$).

B. Averaged values (n = 7) of peak to peak amplitude MEPs recorded during observation of lifting the “Heavy” and the “Light” object. Whiskers indicate standard errors. Asterisks indicate significant differences ($p<.05$). Vertical bars denote ± standard errors.

Figure 4.

Results of experiment 2.

A. Averaged values (n = 5) of muscle activity (EMG) recorded during the **execution** of object lifting with a whole hand grip (expressed as a percentage of the subjects' maximal voluntary contraction (MVC)).

B. Averaged values (n = 12) of peak to peak amplitude MEPs recorded during the **observation** of object lifting with a whole hand grip. Lifting of three different weights was observed: an empty (0 kg), a half full (1 kg) and a full (2 kg) bottle. MEPs evoked from the OP and FCR muscle are presented. Whiskers indicate standard errors. Asterisks indicate significant differences ***p<.001 **p<.01; *p<.05. Vertical bars denote ± standard errors.

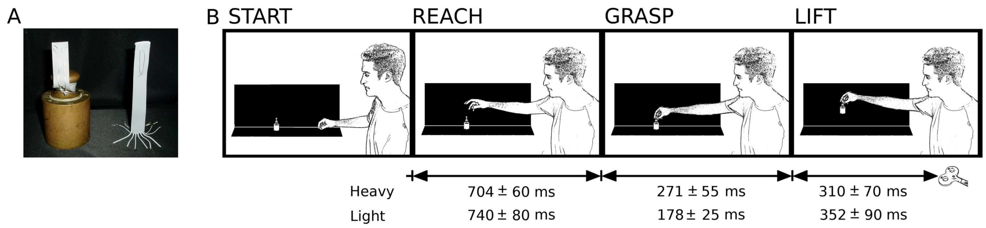

Figure 1.
Experimental stimuli of experiment 1.
A. Picture of the two objects grasped and lifted in front of the subject. The 500g "Heavy" object (left) was a typical brass balance weight. The 10g "Light" (right) object was a piece of ribbon cable. Both objects were grasped using the same grip hand shape.
B. Illustration of events sequence during observation of the reach-grasp-lift action executed upon the heavy object: the actor started hand pronated, then reached to the object, grasped it with precision grip, lifted it and held it over the table during 1s. A TMS pulse was delivered during the lifting phase. Time-line provides the averaged intervals (mean ± std, n=7) between the main task events (button release, hand-object contact, lift onset, TMS pulse) for action upon both the heavy and the light object.

839x214mm (93 x 92 DPI)

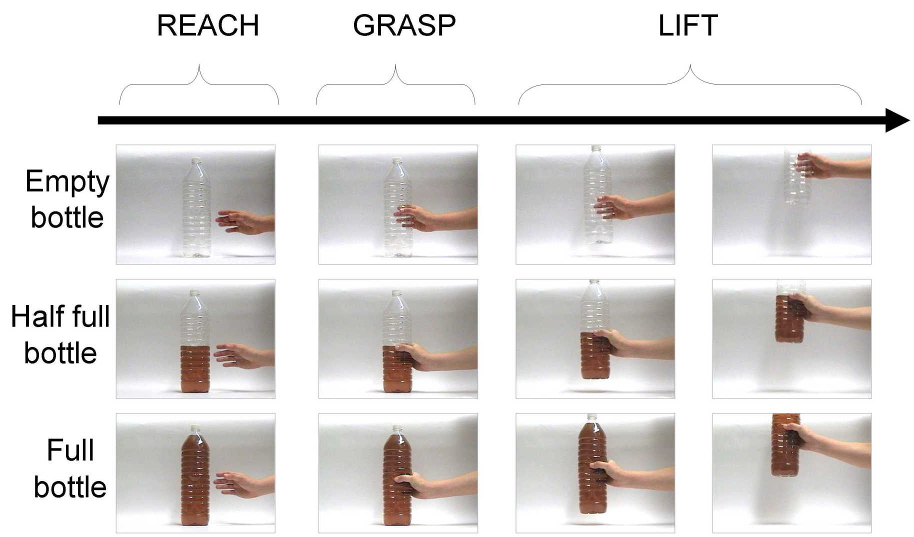

Figure 2.
Experimental stimuli of experiment 2.
The experimental video clips showed a reach-grasp-lift action of a plastic drinking bottle with three different weights, i.e. an empty (0 kg), a half full (1 kg) and a full (2 kg) bottle. The actor entered the scene from the right side, reached to the object, grasped it with a whole hand grip and lifted it out of the scene in the vertical plane. TMS pulses were delivered at random time points during the bottle lifting phase.

639x399mm (122 x 122 DPI)

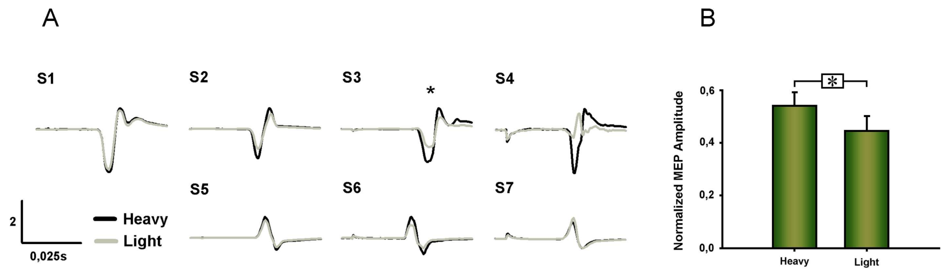

Figure 3.
Results of experiment 1
A. Individual mean MEP traces for the "Heavy" (black) and "Light" (grey) experimental conditions. Asterisks indicate significant differences (p<0.05).
B. Averaged values (n = 7) of peak to peak amplitude MEPs recorded during observation of lifting the "Heavy" and the "Light" object. Whiskers indicate standard errors. Asterisks indicate significant differences (p<.05). Vertical bars denote ± standard errors.

975x296mm (80 x 80 DPI)

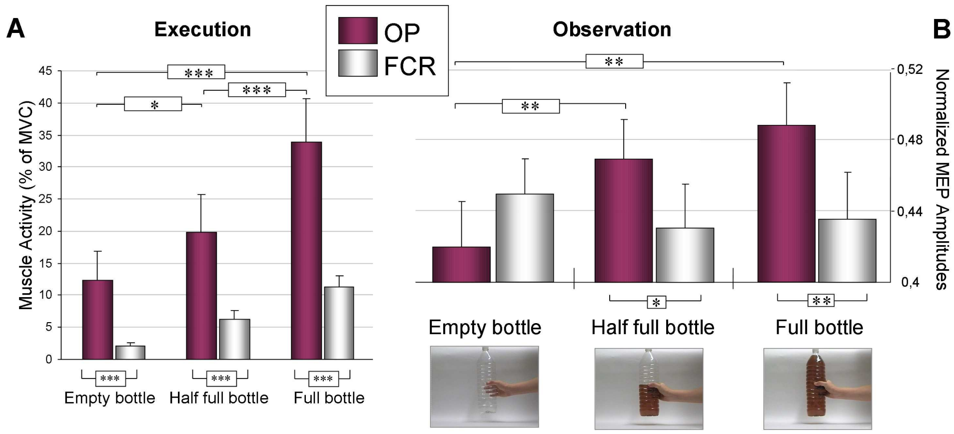

Figure 4.
Results of experiment 2.
A. Averaged values (n = 5) of muscle activity (EMG) recorded during the execution of object lifting with a whole hand grip (expressed as a percentage of the subjects' maximal voluntary contraction (MVC)).
B. Averaged values (n = 12) of peak to peak amplitude MEPs recorded during the observation of object lifting with a whole hand grip. Lifting of three different weights was observed: an empty (0 kg), a half full (1 kg) and a full (2 kg) bottle. MEPs evoked from the OP and FCR muscle are presented. Whiskers indicate standard errors. Asterisks indicate significant differences ***p<.001 **p<.01; *p<.05. Vertical bars denote ± standard errors.

951x449mm (82 x 82 DPI)

FENS Forum 2008

- For posters, morning sessions: 9:30-13:30; afternoon sessions 13:30- 17:30.
  Authors are expected to be in attendance at their posters at the time indicated.
- For other sessions, time indicates the beginning and end of the sessions.

First author: Senot, Patrice (poster)

Poster board D26 - Mon 14/07/2008, 16:00 - Hall 1
Session 123 - Visuo-motor
Abstract n° 123.26
Publication ref.: *FENS Abstr., vol.4, 123.26, 2008*

Authors Senot P. (1, 3), D'Ausilio A. (1), Franca M. (1), Caselli L. (1), Craighero L. (1) & Fadiga L. (1, 2)

Addresses *(1) Univ of Ferrara, Ferrara, Italy; (2) The Italian Inst of Technology, Genova, Italy; (3) Univ Paris 5, Paris, France*

Title Implicit coding of observed kinematics: the case of lifting object with different weights.

Text Mirror neurons in the monkey premotor cortex have been shown to respond similarly for grasping action observation and execution (di Pellegrino et al., 1992). In humans similar mechanisms have been demonstrated with different techniques (Rizzolatti, Craighero, 2004). Primary motor cortex was shown to be facilitated by observation of goal-directed and non-goal-directed actions and this effect was specific for muscles involved in the observed movement (Fadiga et al., 1995). We studied if this facilitation is scaled according to the amount of muscle activity and resulting different kinematics in the observed movement. Moreover we verified if explicit contextual cues could modulate such effect.

Motor Evoked Potentials (MEP) elicited by TMS stimulation of the First Dorsal Interosseous (FDI) muscle representation were measured during the observation of reach-grasp-lift actions upon 6 different objects: 1) transparent empty bottle (VisLight); 2) transparent full bottle (VisHeavy); 3) opaque empty bottle (HidLight); 4) opaque full bottle (HidHeavy); 5) opaque full bottle labeled light (LabLight); 6) opaque full bottle labeled heavy (LabHeavy). Light objects were 50 and heavy were 500g. This difference translated into clear different pattern of muscle contraction and kinematics. TMS was applied when this difference was found to be maximal, in a 100ms window after the beginning of lifting. Condition 1 and 2 afforded full knowledge of weight difference and kinematics information, 3 and 4 only kinematics, 5 and 6 no kinematics but cognitive cues (labels).

We found significant difference in MEPs amplitude in the 1-2 and 3-4 comparisons but no difference in 5-6, 1-3 and 2-4. Results show that the motor cortex does scale for the amount of muscle activity present in the observed action by analysing movement kinematics. Moreover all subjects reported the presence of only 5 objects (they recognised only one opaque) arguing for an implicit processing carried out by the motor system.

Theme D - Sensory and motor systems
Visuomotor processing

Close window

Learning the Nonlinear Multivariate Dynamics of Motion of Robotic Manipulators

E. Gribovskaya , S. M. Khansari Zadeh , Aude Billard

Abstract—**Motion imitation requires reproduction of a dynamical signature of a movement, i.e. a robot should be able to encode and reproduce a particular path together with a specific velocity and/or an acceleration profile. Furthermore, a human provides only few demonstrations, that cannot cover all possible contexts in which the robot will need to reproduce the motion autonomously. Therefore, the encoding should be able to efficiently generalize knowledge by generating similar motions in unseen context.**

This work follows a recent trend in Programming by Demonstration in which the *dynamics* of the motion is learned. We present an algorithm to estimate multivariate robot motions through a Mixture of Gaussians.

The strengths of the proposed encoding are three-fold: i) it allows to generalize a motion to unseen context; ii) it provides fast on-line replanning of the motion in the face of spatio-temporal perturbations; iii) it may embed different types of dynamics, governed by different attractors.

The generality of the method to estimate arbitrary non-linear motion dynamics is demonstrated by accurately estimating a set of known non-linear dynamical systems. The platform-independency and real-time performance of the method are further validated to learn the non-linear dynamics of motion in an industrial six degree of freedom robotic arm and in a four degree of freedom humanoid arm.

Index Terms—**Non-Linear Autonomous Dynamical Systems Robot Programming by Demonstration Learning by Imitation Gaussian Mixture Model and Regression**

I. Introduction

The versatility of tasks that modern robots should accomplish has forced researchers to consider alternative methods for control. Designing task- and robot-specific controllers seems nowadays a time-consuming and ineffective solution, and preference gradually changes in favor of flexible and generic control methods that can adapt to various tasks and robots' geometries. If, in addition, the robot is expected to operate in the vicinity of or in collaboration with unskilled human users, control must be both intuitive and flexible to ensure safe and easy operability by the human.

Programming by Demonstration (PbD) has appeared as one way to respond to this growing need for intuitive control methods [Billard et al., 2008]. PbD designs user-friendly methods by which a human teaches a robot how to accomplish a given task, simply by demonstrating this task. One of the requirements for such a teaching method to be effective is that the number of training examples should remain small (one considers between five and ten examples to be a bearable number for the trainer). Consequently, PbD either relies on prior knowledge to speed up learning, or results in a partial representation of the task which can be refined later.

PbD operates at different levels of the task representation: from copying low-level features of the motion [Sternad and Schaal, 1999, Ude et al., 2004, Calinon and Billard, 2008, Nguyen-Tuong et al., 2008, Schaal et al., 2003] to inferring the user's intention using a symbolic representation [Demiris and B.Khadhouri, 2006, Zollner et al., 2004]. In this paper, we focus on a low-level representation of motions, therefore we further review work related to this direction of PbD. Low-level representations should determine the encoding of the demonstrated trajectories of motion so that they can be easily modulated to enable re-use of the skill in novel contexts. An overview of requirements for effective movement encoding has been summarized in [Ijspeert et al., 2001].

Most relevant to the present paper are the notions of *compactness* and *reusability* of the representation, i.e. the encoding should be easily transferrable to related tasks, and the notion of *robustness* to perturbations, i.e. an ability of an encoding to ensure that a motion may be quickly adapted to perturbation and changes in a dynamic environment.

Dynamical Systems (DS) provide an effective and elegant means of encoding motions, that fulfills the above three criteria. DS encode trajectories through a time-independent function that defines the temporal evolution of the motion. Generalization of the motion to an unobserved part of the space results immediately from the application of the function to the new set of input variables.

In this paper, we consider the problem of estimating a *time-independent* model of motion through a set of first order non-linear multivariate dynamical systems. We exploit the strength of parametric statistical techniques to learn correlations across the variables of the system and show that this technique allows the determination of a coarse representation of the dynamics. We demonstrate advantages of such an approach as an alternative to the *time-dependent* methods, by ensuring robustness to external spatio-temporal perturbations through on-line adaptation of the motion.

This paper is divided as follows. Section II reviews related work on motion learning and estimation of dynamical systems. Section III-A starts with a formalization of the problem at hand and the particular approach of this work. This is followed by a technical description of the modeling approach: Section III-B introduces the learning approach to estimate the dynamics, while Section III-C presents an iterative algorithm to improve stability of the learned dynamics. Finally, in Section IV, we validate the method by estimating the motion dynamics from trajectories generated with given dynamical laws; in this way we may systematically verify approximation qualities of the method. We, further, show how the same framework can be used to learn the dynamics of motion of a 4 degree of freedom humanoid robot arm and a 6 degree of freedom industrial arm. The legend used in graphs throughout the paper is summarized in Figure 1. The glossary is in Table I.

II. Related Work

To better delineate this paper's particular contribution to both machine learning and robotics, we focus our review on two major themes. First, to situate the dynamical systems approach taken in our work, we make a brief historical tour of the large volume of literature on modeling robot motion, contrasting time-dependent and time-independent representations. We then turn to the problem of estimating arbitrary dynamical systems and introduce the particular statistical technique used here. We briefly summarize the broad division across parametric and non-parametric statistical methods, and situate our choice of parametric method in this context.

A. Motion Learning

A core issue within robot control is ensuring that, if perturbed, the robot's motion can be rapidly and on-the-fly recomputed to ensure that the robot ultimately accomplishes the task at hand. Perturbations may lead the robot to either depart from its original trajectory (e.g. when slipping or hitting an object) or be delayed (e.g. when slowed down because of friction in the gears). In the rest of this paper, we will refer to the former type of perturbations as *spatial* perturbations and to the latter as *temporal* perturbations.

The vast majority of work on motion learning has addressed essentially the problem of being robust to spatial perturbation. Very little work has been yet done on handling temporal perturbations, which is core to the model we develop here. Next, we review these different approaches.

B. Time-dependent Modeling Approaches

Traditional means of encoding trajectories are based on spline decomposition after averaging across training trajectories [Hwang et al., 2003, Andersson, 1989, Yamane et al., 2004, Aleotti et al., 2005]. Spline decomposition remains a powerful tool for quick trajectory formation. It is, however, heavily dependent on a heuristic for segmenting and aligning the trajectories. Furthermore, spline representation, not being statistically-based, may have difficulties in coping with noise in data that is inherent in the robotic application.

Non-linear regression techniques were proposed as a statistical alternative to spline-based representation [Calinon et al., 2007, Schaal and Atkeson, 1998, 1994, D. et al., 2008]. These methods allow the systematical treatment of uncertainty by assuming the noise in data and, therefore, by estimating actual trajectories as a set of random variables with learned parameters.

However, similarly to spline-based approaches, regression techniques depend on an explicit time-indexing and virtually operate in "open-loop". The lack of any kind of feedback makes regressions sensitive to both temporal and spatial perturbations. To compensate for this, one needs to introduce an external mechanism to track potential deviations from the desired trajectory during reproduction. Adaptation to deviations then relies on a heuristic to re-index the new trajectory in time or extrapolate in space. Such re-indexing or extrapolation often comes at the cost of deviating importantly from the desired velocity and acceleration profile, making the motion look "unnatural". Furthermore, finding a good heuristic is highly task-dependent and becomes particularly not-intuitive in multidimensional spaces [Schaal et al., 2003].

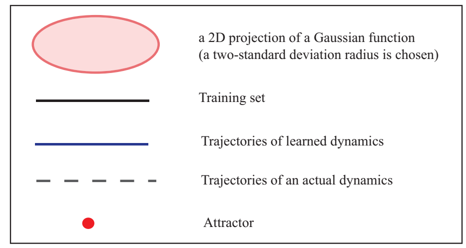

Fig. 1. Legend for the Figures in the paper.

Time-independent models, such as autonomous dynamical systems (to which we will further refer to as DS), were recently advocated as an alternative to the above approaches[1]. Models based on DS are advantageous in that they do not depend on an explicit time-indexing and thus provide a closed-loop controller, while being able to model arbitrary non-linear dynamics. Removing the explicit time-dependency comes at a cost, as it re-introduced an old problem, namely the need to consider stability of the control policy.

Next, we review current approaches to DS modeling of robot motion and point out the limitations of these methods. For a detailed discussion on advantages and disadvantages of dynamical systems encoding of motion, see also [Ijspeert et al., 2001, Schaal et al., 2003, 2001, Schoner and Santos, 2001].

C. Dynamical Systems Modeling of Motion

A number of recent approaches in PbD, including our prior work, investigate the use of dynamical systems for modeling robot motions [Ijspeert et al., 2001, Righetti et al., 2006, Dixon and Khosla, 2004, Ijspeert and Crespi, 2007, Hersch et al., 2008]. While [Dixon and Khosla, 2004] focuses on fitting the parameters of a first-order linear dynamical system into training data, the other above works tackle a problem of modulating a predefined linear dynamics with a non-linear estimate of a trajectory [Hersch et al., 2008] or a velocity profile [Ijspeert et al., 2001, Righetti et al., 2006]. The authors choose an uni-variate spring and damper system as an underlying linear dynamics. In such a way, they avoid an issue of stability of approximation that may occur if one learns an actual dynamics from data. However, this solution comes with its drawbacks: (1) uni-variate encoding discards information about correlation between degrees of freedom, that may be crucial for faithful reproduction (see Figure 23 for

[1] DS formulation embeds the time-dependency of a system in the mathematical formulation of the problem by using time derivatives of the state variables.

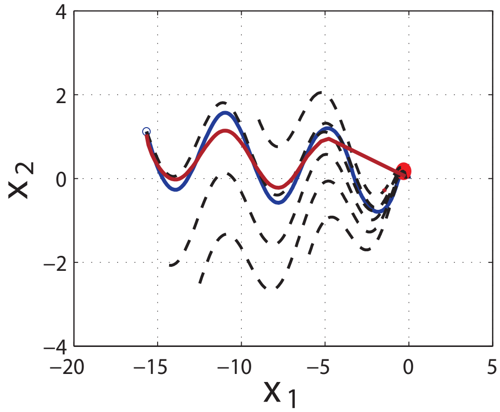

Fig. 2. A two-dimensional theoretical dynamical system is estimated from five training samples (dotted lines). using Dynamical Motion Primitives [Ijspeert et al., 2002](red solid line) and the method proposed in this paper (blue solid line). As approximation with Dynamical Motion Primitives requires combining statistical results with a predefined dynamics, it may deform an actual path to follow, as it can be seen on the graph.

illustration of the uni-variate encoding problem). (2) Coupling of the output of a predefined linear DS with a regression estimate makes the overall system dependent on the temporal synchronization between the two signals and thus in effect time-dependent (see Table VI for a formal comparison between the proposed approach and the work [Ijspeert et al., 2001]). To handle temporal perturbations, one would need a heuristic to maintain the synchronization. This would, however, no longer guarantee that the overall system is globally asymptotically stable. (3) By ensuring that the stable DS takes precedence over the estimate when coming close to the attractor or after a given time period, one can show global stability of the complete estimate [Ijspeert et al., 2002]. In effect, the global dynamics of motion is increasingly dominated by the stable linear dynamical system, hence leading the motion to progressively depart from the learned dynamics. This effect is illustrated in Figure 2, where we see that the trajectory is distorted as the system approaches the target. To ensure that the modulation still influences the dynamics of the motion when approaching the target, the method relies on using a large number of Gaussians spread across the data points.

In this paper, we develop an iterative procedure to learn a statistical estimate of an arbitrary multivariate autonomous dynamical system. We discuss the problem of stability of a learned estimate and propose an empirical procedure to verify stability and the region of applicability of the estimate. This relieves us from the need of using another a priori stable dynamical system and ensures robustness against spatial and temporal perturbations.

D. *Estimating a Dynamical System*

Data-driven methods for estimating dynamical systems consider multivariate input-output data as instances of a dynamical system and seek an estimate of the model that relates best these pairs of datapoints. Building a local approximation of the dynamics has been first reviewed within the time series analysis [Priestley, 1980, Chamroukhi et al., 2009, Ljung, 2004]. These works consider solely uni-dimensional data with a major motivation of predicting time series.

Analysis of dynamics has gradually shifted to state-space representation as it allows a representation of more sophisticated phenomena [Aoki, 1990, Crutchfield J.P., 1987]. The vast majority of these works focus on estimating *linear* dynamics [Dixon and Khosla, 2004, Ryoung K. Lim, 1998], a restrictive assumption for robotic applications. Recently, with the growing interest in chaos theory, more developed approaches have been proposed that allow approximation of complex dynamics [Crutchfield J.P., 1987, Wang et al., 2008]. While, several optimistic results in simulations have been presented [Carroll, 2007, Xie and Leung, 2005], their applicability to practical tasks with a small number of observed data containing noise remains to be verified.

The major body of numerical approaches of non-linear dynamical systems perform function approximation using different orthogonal polynomials (Chebyshev polynomials, B-splines [Lee, 1986], Radial Basis Functions [Buhmann, 2003] (RBFs)). Recently, many works have addressed the approximating properties of RBFs [Tomohisa et al., 2008, Travis et al., 2009, Wei and Amari, 2008]. RBFs have been proved to form universal approximators of any function on a compact set [Park and Sandberg, 1991]: any level of precision of the approximation may be achieved by considering an exhaustive number of basis functions; however, the quality of the approximation heavily depends on tuning a considerable amount of parameters. Thus, the problem of determining a tuning procedures optimum according to different criteria is a recurrent subject in the domain [Buhmann, 2003]. Furthermore, as the approximation with RBFs falls naturally into the category of non-parametrical methods discussed next, they suffer from the same types of limitations: RBFs better suits for approximation of uni-variate signals and quality of approximation rapidly deteriorates with an increase in the number of dimensions.

E. *Statistical Encoding*

Classically, the whole body of statical methods can be broadly divided into *parametric* and *non-parametric* approaches.

Non-parametric methods used in robot motion estimation include k-nearest neighbors [Moore, 1990], Gaussian Processes [Deisenroth et al., 2009, Nguyen-Tuong et al., 2008], Locally Weighted Regression, [Hardle, 1991, Muller, 1988, Schaal and Atkeson, 1998, 1994] and a combination of these [Nguyen-Tuong et al., 2008]. Non-parametric methods are advantageous over parametric methods as they make little assumptions about the form of the underlying distribution function to estimate. Moreover, due to the local nature of their estimate, non-parametric methods are well suited for accurate data fitting in low-dimensional spaces [Schaal and Atkeson, 1998, 1994]. Initially proposed for uni-dimensional problems, the above non-parametrical methods suffer from the curse of dimensionality [Bellman, 1957]: sparsity of training data in high-dimensional spaces makes accurate estimation of parameters almost impossible. Parametric methods, in contrast, are better suited to model a multivariate dataset. They, however, rely on heuristics to choose the underlying parameters efficiently.

The Gaussian Mixture Models (GMMs) and based on them Gaussian Mixture Regression (GMR) are parametric methods. They are thus better suited for regression on multi-dimensional data [Sung, 2004]. Learning with GMM is classically done using Expectation-Maximization (EM), the iterative algorithm that optimizes the likelihood of the mixture of Gaussians over the data. Optimal performance relies, however, on choosing the number of Gaussians and on the stopping criterion of EM (see [McLahlan and Peel, 2000] for a review). While several methods have been proposed to automatically estimate these two parameters, with the Bayesian Information Criterion (BIC[2]) being the most generic, GMM estimation using EM may lead to suboptimal results and remain very sensitive to the initialization conditions. Here, we show that, for both our problem at hand and in practice, these known limitations are not an impediment and that an iterative method for choosing the number of Gaussians leads to good performance. Most importantly, we show that the method converges quickly and relies on very few parameters in comparison to parametric methods.

III. Method

A. Problem Statement

Consider that the state[3] of our robotic system can be unambiguously described by a variable ξ and that the workspace of the robot forms a sub-space X in $\mathbb{R}^N$.

Consider further that the state of our robotic system is governed by an *Autonomous Dynamical System* $\langle \mathcal{X}, f, \mathcal{T} \rangle$ (as per Definition 1-2, Table I). Then, for all starting locations $\xi_0 \in X$, the temporal evolution of our robotic system is *uniquely determined* by the *state transition map* (Definition 2, Table I) $f(t, t_0, \xi_0) = \xi(t)$, $\forall \xi_0, \xi \in X$.

Let us further assume that the state transition map f is a non-linear continuous and continuously differentiable function and that the system is driven by a first order differential equation[4] with a single equilibrium point $\bar{\xi}$, such that:

$$\forall t \in T = [t_0; \infty]; [\xi; \dot{\xi}] \in X \subset \mathbb{R}^N \quad (1)$$

$$\dot{\xi}(t) = f(\xi(t)) \quad (2)$$

$$\dot{\bar{\xi}} = f(\bar{\xi}) = 0. \quad (3)$$

Let the set of M N-dimensional demonstrated datapoints $\{\xi^i, \dot{\xi}^i\}_{i=1}^M$ be instances of the above motion model. The problem consists then of building an estimate $\hat{f}$ of f based on

TABLE I
GLOSSARY OF DEFINITIONS

Definition 1: The *state-space* $X \subset \mathbb{R}^N$ includes all possible instantiations of ξ, such that $\xi(t) \in X$ at each time step $t \in T = \mathbb{R}^+ = [0; \infty]$.

Definition 2: A *dynamical system* is the tuple $\langle \mathcal{X}, f, \mathcal{T} \rangle$, with $f : t \rightarrow f^t$ a continuous map of X onto itself.

Definition 3: A dynamical system is *differentiable* if $\exists f : T \times X \rightarrow X$ such that for all $t_0 \in T, \xi_0 \in X$ the problem:

$$\dot{\xi} = f(t, \xi(t)), \;\; t \geq t_0, t \in T$$

$$\xi(t_0) = \xi_0$$

has a unique solution.

A dynamical system governed by a time-independent transition map with $f(t, \xi(t)) \triangleq f(\xi(t))$ is an *Autonomous Dynamical System*.

Definition 4. An *equilibrium state* $\bar{\xi} \in X$ of a dynamical system is such that

$$f(t, t_0, \bar{\xi}) = \bar{\xi}.$$

Definition 5. An equilibrium state $\bar{\xi} \in X$ is *stable* if $\exists \epsilon > 0$ and $\delta = \delta(\epsilon)$ such that

$$\forall \xi_0 \in B(\bar{\xi}, \delta) \Rightarrow f(\xi_0) \in B(\bar{\xi}, \epsilon),$$

$B(\bar{\xi}, \delta) \subset X$ is a hypersphere centered at $\bar{\xi}$ with radius δ. $\bar{\xi}$ is an *attractor* of f.

Definition 5. An *attractive state* is an equilibrium state $\bar{\xi}$ of a local flow, if there exists $\rho > 0$ such that:

$$\forall \xi_0 \in B(\bar{\xi}, \rho) \Rightarrow \lim_{t \to \infty} f(\xi_0) = \bar{\xi}.$$

$B(\bar{\xi}, \delta) \subset X$ is a hypersphere centered at $\bar{\xi}$ with radius δ. $\bar{\xi}$ is an *attractor* of f.

Definition 6. An equilibrium point $\bar{\xi}$ is *asymptotically stable* if it is both stable and attractive.

Definition 7. A set $\Delta \subset X$ is a *Region of Attraction (or Basin of Attraction)* of an equilibrium $\bar{\xi}$ if:

$$\Delta(\bar{\xi}) = \{\xi_0 \in X; \lim_{t \to \infty} f(\xi_0) = \bar{\xi}\}$$

See Figure 23-II for illustration.

Definition 8. A dynamical system is *globally asymptotically stable* at the equilibrium $\bar{\xi}$ if $\bar{\xi}$ is an asymptotically stable attractor and $\Delta \equiv X$.

the set of demonstrations. To this end, we will approximate the function in a subregion[5] $C \subset X$, so that:

$$\hat{f} : C \rightarrow C \quad (4)$$
$$\hat{f}(\xi(t)) \approxeq f(\xi(t)), \forall \xi \in C.$$

C is further referred to as the *region of applicability* of a learned dynamics.

Without loss of generality, we can transfer the attractor to the origin[6], so that $\bar{\xi} = 0 \in C \subset X$ is now the equilibrium point of f and by extension of its estimate $\hat{f}$, i.e. $\hat{f}(0) = f(0) = 0$. If C is contained within the *region of attraction* Δ of $\bar{\xi}$ (see Definition 7, Table I), then the estimate $\hat{f}$ is asymptotically stable at $\bar{\xi}$ in C and any motion initiated from $\xi(t_0) \in C$ will asymptotically converge to the target $\bar{\xi}$.

[2]BIC introduces a penalty term for increasing the number of parameters in the model over the resulting improvement in the modeling performance.

[3]The state of a dynamical system represents the *minimum amount of information* required to describe the effect of past history on the future development of this system [Hinrichsen D., 2000].

[4]Considering solely *first order* dynamical systems is not restrictive to learning only first order relationships between trajectory and velocity, as one can always convert dynamics of an arbitrary order into a canonical system of first order ODEs.

[5]Estimating the dynamics in the whole state-space X would be practically infeasible due to the excessive number of demonstrations that this would require.

[6]To simplify the notation, we keep the same notation for the domains C and X after translation at the origin.

B. Approximating the Dynamics with Gaussian Mixture Regression

To construct $\hat{f}$ from the set of demonstrated trajectories, we follow a statistical approach and define $\hat{f}$ as a non-linear combination of a finite set of Gaussian kernels, using *Gaussian Mixture Models* (GMM).

GMMs define a joint probability distribution function $\mathcal{P}(\xi^i, \dot{\xi}^i)$ over training set of demonstrated trajectories $\{\xi^i, \dot{\xi}^i\}$, $i = 1..M$, M is the number of demonstrations, as a mixture of a finite set of K Gaussians $G^1..G^K$ (with μ^k and Σ^k being the mean value and covariance matrix of a Gaussian G^k):

$$\mathcal{P}(\xi^i, \dot{\xi}^i) = \frac{1}{K} \sum_{k=1}^{K} G^k(\xi^i, \dot{\xi}^i; \mu^k, \Sigma^k) \tag{5}$$

and

$$\mu^k = [\mu^k_\xi \ ; \ \mu^k_{\dot{\xi}}] \ \text{ and } \ \Sigma^k = \begin{pmatrix} \Sigma^k_\xi & \Sigma^k_{\xi\dot{\xi}} \\ \Sigma^k_{\dot{\xi}\xi} & \Sigma^k_{\dot{\xi}} \end{pmatrix} \tag{6}$$

Where each Gaussian probability distribution G^k is given by:

$$G^k(\xi^i_t, \dot{\xi}^i_t; \mu^k, \Sigma^k) = \tag{7}$$
$$\frac{1}{\sqrt{(2\pi)^{2d}|\Sigma^k|}} e^{-\frac{1}{2}(([\xi^i_t, \dot{\xi}^i_t] - \mu^k)^T (\Sigma^k)^{-1} ([\xi^i_t, \dot{\xi}^i_t] - \mu^k))}.$$

The model is initialized using the *k-means* clustering algorithm starting from a uniform mesh and refined iteratively through Expectation-Maximization (EM) [Dempster et al., 1977].

To generate a new trajectory from learned GMMs, one can then sample from the probability distribution function given by Eq.5. This process is called Gaussian Mixture Regression (GMR).

Taking the posterior mean estimate of $\mathcal{P}(\dot{\xi}|\xi)$, the estimate of our function $\hat{\dot{\xi}} = \hat{f}(\xi)$ can then be expressed as a non-linear sum of linear dynamical systems, given by:

$$\hat{\dot{\xi}} = \sum_{k=1}^{K} h_k(\xi)(A_k \xi + B_k), \tag{9}$$

where $A_k = \Sigma^k_{\dot{\xi}\xi}(\Sigma^k_\xi)^{-1}$, $B_k = \mu^k_{\dot{\xi}} - A_k \mu^k_\xi$, $h_k(\xi) = \frac{\mathcal{P}(\xi; \mu^k_\xi, \Sigma^k_\xi)}{\sum_{k=1}^{K} \mathcal{P}(\xi; \mu^k_\xi, \Sigma^k_\xi)}$, $h_k(\xi) > 0$, and $\sum_{k=1}^{K} h_k(\xi) = 1$.

Such a rewriting will prove useful when studying the stability of the estimate, as will be discussed in Section III-C.

A geometric illustration of the GMR inference in the case of single Gaussian is presented in Figure 3 and the GMR procedure is summarized in Table II. Figure 4 further illustrates the encoding process from GMM to GMR for a non-linear dynamical system with a single attractor.

TABLE II
GAUSSIAN MIXTURE REGRESSION

Let us assume that we can for each input datapoint $\xi^{\mathcal{I}}$ match an output datapoint $\xi^{\mathcal{O}}$, the joint probability of input and output data is then modeled using Gaussian Mixtures. The probability that a datapoint $\eta = [\xi^{\mathcal{O}}; \xi^{\mathcal{I}}]$ belongs to the GMM is defined by

$$\mathcal{P}(\eta) = \sum_{k=1}^{K} \pi_k \ \mathcal{N}(\eta; \mu_k, \Sigma_k) =$$
$$= \sum_{k=1}^{K} \pi_k \ \frac{1}{\sqrt{(2\pi)^D |\Sigma_k|}} \ e^{-\frac{1}{2}\left((\eta - \mu_k)^\top \Sigma_k^{-1} (\eta - \mu_k)\right)}$$

where π_k are prior probabilities and $\mathcal{N}(\mu_k, \Sigma_k)$ are Gaussian distributions defined by centers μ_k and covariance matrices Σ_k, where input and outputs components are represented separately as

$$\mu_k = \begin{bmatrix} \mu_k^{\mathcal{I}} \\ \mu_k^{\mathcal{O}} \end{bmatrix}, \quad \Sigma_k = \begin{bmatrix} \Sigma_k^{\mathcal{I}} & \Sigma_k^{\mathcal{IO}} \\ \Sigma_k^{\mathcal{OI}} & \Sigma_k^{\mathcal{O}} \end{bmatrix}.$$

Gaussian Mixture Regression allows to compute for a given input variable $\xi^{\mathcal{I}}$ and a given component k, the expected distribution of $\xi^{\mathcal{O}}$ as:

$$\mathcal{P}(\xi^{\mathcal{O}}|\xi^{\mathcal{I}}, k) \sim \mathcal{N}(\hat{\eta}_k, \hat{\Sigma}_k), \text{ where } \quad \begin{aligned} \hat{\eta}_k &= \mu_k^{\mathcal{O}} + \Sigma_k^{\mathcal{OI}}(\Sigma_k^{\mathcal{I}})^{-1}(\xi^{\mathcal{I}} - \mu_k^{\mathcal{I}}), \\ \hat{\Sigma}_k &= \Sigma_k^{\mathcal{O}} - \Sigma_k^{\mathcal{OI}}(\Sigma_k^{\mathcal{I}})^{-1}\Sigma_k^{\mathcal{IO}}. \end{aligned}$$

where $h_k = \mathcal{P}(k|\xi^{\mathcal{I}})$ is the probability of the component k to be responsible for $\xi^{\mathcal{I}}$

$$h_k = \frac{\mathcal{P}(k)\mathcal{P}(\xi^{\mathcal{I}}|k)}{\sum_{i=1}^{K} \mathcal{P}(i)\mathcal{P}(\xi^{\mathcal{I}}|i)} = \frac{\pi_k \ \mathcal{N}(\xi^{\mathcal{I}}; \mu_k^{\mathcal{I}}, \Sigma_k^{\mathcal{I}})}{\sum_{i=1}^{K} \pi_i \ \mathcal{N}(\xi^{\mathcal{I}}; \mu_i^{\mathcal{I}}, \Sigma_i^{\mathcal{I}})}.$$

Alternatively, by using the linear transformation property of Gaussian distributions, the conditional expectation of $\xi^{\mathcal{O}}$ given $\xi^{\mathcal{I}}$ can be defined approximately defined by a single normal distribution with the parameters:

$$\hat{\mu} = \sum_{k=1}^{K} h_k \ \hat{\mu}_k \ , \quad \hat{\Sigma} = \sum_{k=1}^{K} h_k^2 \ \hat{\Sigma}_k. \tag{8}$$

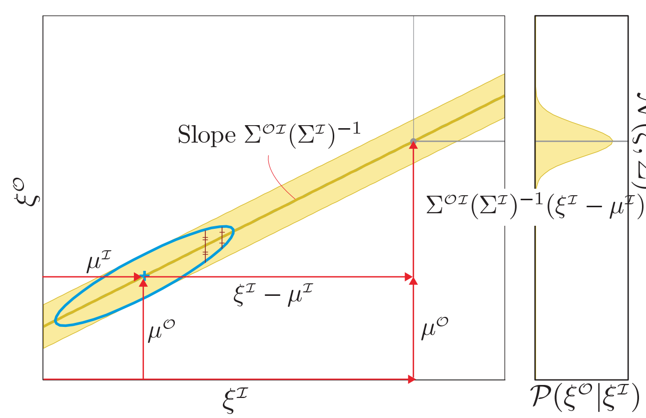

Fig. 3. The geometric illustration of Gaussian Mixture Regression inference (see also Table II). GMR approximates our dynamical systems through a non-linear weighted sum of local linear models: each regression matrix $A_k = \Sigma_k^{\mathcal{OI}}(\Sigma_k^{\mathcal{I}})^{-1}$ defines coefficients of the local linear fit. Here, we display the effect of fitting with a single Gaussian a pair of input and output signals $\xi_i^O \in R^M$, $\xi_i^I \in R^P$ respectively. The projection of the regression signal to the subspace spanned by $\{\xi_i^{O,m}, \xi_i^{I,p}\}$ is a line with a slope given by the elements A_k^{mp} of the regression matrix (i.e., $\xi_i^{O,m} = A_k^{mp}\xi_i^{I,p}$). The mixture of covariance matrix in GMM defines a probabilistic envelope around the regression signal. Thus, to each input ξ_i^I is associated a probability distribution function for output $\mathcal{P}(\xi_i^O|\xi_i^I)$, with mean ξ_i^O. In the present work, we exploit the envelope to determine the boundaries for our generalized inverse kinematics solution when the solution is not exactly the regression signal ξ_i^O.

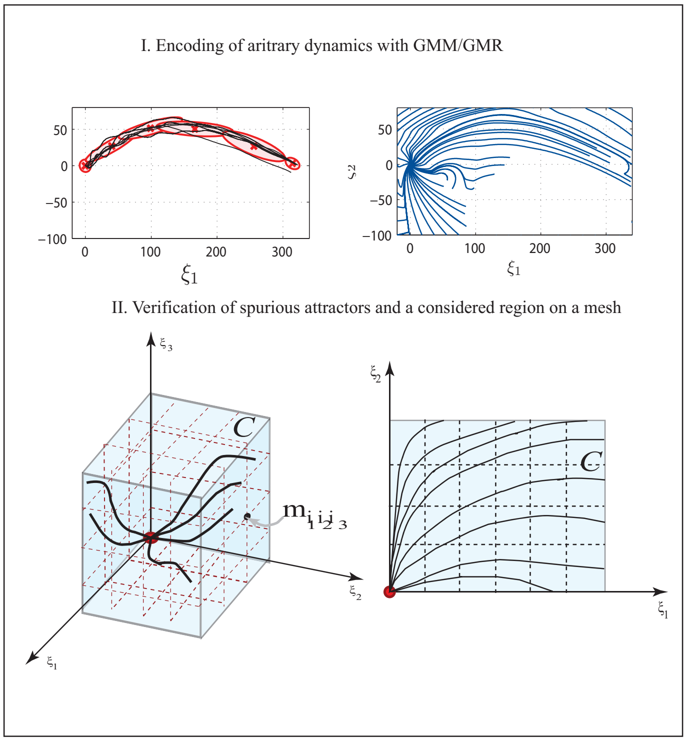

Fig. 4. I. Illustration of a GMM/GMR encoding of an arbitrary dynamics. *Top left*: Two-dimensional projection of the data with superimposed the Gaussian Mixture envelope. Top right: All trajectories regenerated using Gaussian mixture regression when starting from 20 different locations in space converge correctly to the the origin, the attractor of the system. *Bottom left and right*: in blue, the region of applicability C that embeds all demonstrated trajectories. To empirically determine if C is a region of attraction, C is sampled equally and one measures if all trajectories originating from each of sampled point converges correctly to the target.

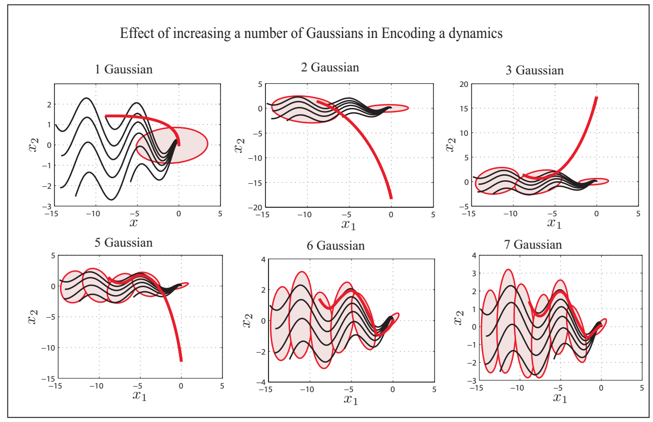

Fig. 5. Improvement in the stability of approximation with the increase the number of Gaussian components

C. Stability Analysis

Stability analysis of *linear* dynamical systems is well-studied subject [Khalil, 1996]: one either constructs a Lyapunov function for the system or analyzes the eigenvalues of the control matrix.

In contrast, there is no unique method to analyze the stability of *non-linear* dynamical systems and theoretical solutions exist only for particular cases. Classically, stability analysis of non-linear dynamical systems is performed in two steps: first, the system is linearized in a neighborhood around the points of interest (the attractors) and their asymptotic stability is verified; second, analysis of the region around the attractors is done to determine the extent of the region of attraction.

Methods to analytically estimate the regions of attraction (see Definition 7, Table I) are often based on the construction of a Lyapunov function gradually expanding its region of validity [Bai et al., 2007, Giesl, 2008, Genesio et al., 1985]. Such a procedure however produces a rather coarse estimation of the region of attraction and may fail to identify regions with non-convex boundaries. Alternative approaches take a geometrical perspective by reversing the flow of motion (by analyzing a dynamical function with a opposite sign) starting from the attractor and finding repellers and boundaries for a region of attraction from the reversed trajectories [Loccufier and Noldus, 2000]. These methods are more accurate but require considerable computation time, a known structure of an attractor's landscape (number of existing attractors and repellers).

Theoretical estimation of the region of attraction in the general case of multivariate non-linear systems is thus still an open problem. In practice, one relies on numerical procedures for evaluating whether a given region of applicability is a region of attraction. Here, we follow such an approach.

We start from the observation that GMR gives us a non-linear weighted sum of linear dynamical systems; see Eq. 9. Stability of the system is governed by the GMR parameters (the matrices A_k, B_k and mixing coefficients h_k), which are learned during training. Since the stability of the learned dynamics depends on the parameters of the training algorithm (Expectation Maximization) in Section III-D) we will show that a modification of the GMM procedure to build the mixture results in an estimate locally stable around the target.

1) Local stability at the origin: Following from the hypothesis that the origin is an attractor of the true control law $\dot{\xi} = f(\xi(t))$, we must ensure that its estimate given by (9) is also stable at the origin. Recall that for a point to be an attractor of the system (see Definition 5, Table I), there must exist a region around it where all trajectories are asymptotically stable.

Let us assume that in the neighborhood of the origin the system is governed solely by the last Kth gaussian [7]. In other words, let us assume there exists a neighborhood of the origin, where for points ξ in this neighborhood all mixing coefficients expect the Kth are zeros: $\exists B(\epsilon)$such that $\forall \xi \in B(\epsilon)\ h_k(\xi) \simeq 0\ \ k = 1..K-1$, where $B(\epsilon)$ is a hypersphere of radius ϵ. In this region, the system governed by Eq.9 reduces then to:

$$\dot{\xi} = A\xi + B \tag{10}$$

with $A = \Sigma_{K,\dot{\xi}\xi}\Sigma^{-1}_{K,\xi}$ and $B = \mu_{K,\dot{\xi}} - A\mu_{K,\xi}$.

The system above, driven by Eq. 10, will be asymptotically stable if the eigenvalues of the matrix A are all strictly negative. For a $m \times m$-dimensional matrix to be negative definite, all its i-th order leading principal minors should be

[7] In practice, as we seek to avoid the over-fitting, the Gaussians are sufficiently set apart, therefore at the origin the influence of all other Gaussians except for the last one becomes numerically zero.

negative if i is odd and positive if i is even; stability, therefore, is guaranteed when the following set of constraints is satisfied:

$$\|A_{\dot{\xi}\xi,[1:i,1:i]}\|(-1)^i < 0 \quad \forall i = 1, ..., m \text{ that is satisfied if} \tag{11}$$

$$(1)\ a_{ii} < 0 \text{ and } (2)\ a_{ij} \ll a_{ii}\ \forall\ i, j = 1, ..., m \text{ and } i \neq j, \tag{12}$$

where $A_{\dot{\xi}\xi} = \{a_{ij}\}_{i,j=1}^N$.

Figure 6 illustrates geometrically the effect of the local stability condition on the dynamics of motion and the form of the Gaussian. When projected on the $\{\xi_i, \dot{\xi}_i\}$ axes, the Gaussian corresponds to an ellipse with the main axis forming a negative slope. This results in a homogenous flow of motion toward the attractor along all dimensions.

For EM to result in such an elongated Gaussian, training data must homogeneously cover the space of motion around the target. This means that one should show the robot how to approach the target by uniformly starting all around the target. In practice, because the training set is finite and gives only a partial coverage of the state space, GMM estimate will be imprecise, resulting in both a shift of the slope of the Gaussian and a shift of the attractor's location, see Figure 6. Additional measures should, thus, be taken to guarantee the convergence to the target, which we describe next.

D. Practical approach to ensuring and analyzing stability

1) Ensuring local stability empirically: To overcome the lack of uniformly distributed training data around the origin in the experiments presented here, we generate additional so-called *synthetic* data by rotating a subset of training data, selected within a small neighborhood, around the origin. In addition, we set the center of the last Gaussian of the GMM at the target, i.e. at the origin ($\mu_{K,\xi} = \mu_{K,\dot{\xi}} = 0$), and do not update this center during training. This procedure is illustrated in Figure 6.

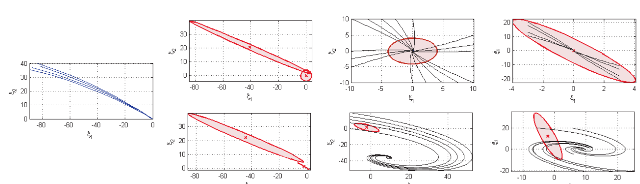

Fig. 6. Influence of the accurate positioning of the last Gaussian at the origin. *Top*: the last gaussian is positioned at the origin through addition of synthetic datapoints, that guarantees asymptotic stability of the system in the neighborhood of the origin, as can be seen from the vector field trajectories (the very right graph). *Bottom*: however, the real data asymptotically converge to the origin (the very left graph), the statistical EM does not automatically position the last Gaussian at the origin, that leads to the convergence to the spurious attractor (the very right graph).

The system driven by the truncated dynamics is given by Eq.(10) and a system generated through this procedure is ensured to be asymptotically stable within a neighborhood around the origin. Next, we describe a procedure by which we can empirically estimate boundaries of the region of applicability C.

TABLE III
MODEL TRAINING

1	Collect a dataset of demonstrations and initialize C (see Section III-D2).
2	Add synthetic data around the target
3	Choose an initial number of GMM components K ($K = 2$ in the experiments reported here)
4	**LOOP** until stable approximation is found
5	Train the joint probability $\mathcal{P}(\xi, \dot{\xi})$ with Expectation Maximization [Dempster et al., 1977]:
6	Verify local stability at the origin Eq. (12)
6	**IF** (the origin is not asymptotically stable) **THEN** increase the number of GMM components $K = K + 1$
	ELSEIF (estimate of C does not include all training trajectories) **OR** ($\exists$ spurious attractors inside the region C) **THEN** add training data **AND** retrain
	END
8	**EN**

2) Determining the region of stability empirically: As mentioned in Section III-A, estimating dynamics in the whole state-space X is impractical. Instead, we will estimate stability locally within a subset $C \subset X$. C includes training data points and lies inside the robot's workspace. Initialization of C is data-driven: size of C along each dimension is defined by the amplitude of the training dataset along this dimension.

After training, the initial guess regarding C should be restimated, to empirically verify that C is a region of attraction of the origin and that it does not include any other attractors . We follow a numerical procedure in which we integrate trajectories forward starting from a uniform mesh defined on the boundaries, and verify that all the trajectories converge toward the origin.

To do this, we construct a mesh M covering boundaries of C: $M(\tau_1..\tau_N) = \{(\xi^1_{i_1}..\xi^N_{i_N}) = (i_1\tau_1..i_N\tau_N),\ i_1 = 1..n_1, ...,\ i_N = 1..n_N\}$, where $\tau_1 = c_1/n_1\ ..\tau_N = c_N/n_N$, $c_1\ ..\ c_N$ – size of each of dimensions of C; $n_1..\ n_N$ – size of the mesh along each of dimensions in $\mathbb{R}^N$ (see Figure 4-II).

We integrate trajectories starting from each node $(\xi^1_{i_1}..\xi^N_{i_N})$ on the mesh M and verify that the velocity is zero only at the origin, thus ensuring that only the origin of the system is an attractor. If this condition is satisfied all trajectories starting inside C will not leave the boundaries, due to the properties of differential equations.

To improve stability, we increment the number of Gaussians K and re-estimate the system using EM. Augmenting the number Gaussians allows a more precise encoding of the dynamics locally along the trajectory; see Figure 5. Since instabilities result often in the motion exiting the desired trajectory (e.g. if there are sharp turns in the trajectory that have been poorly approximated by the mixture), increasing the granularity of the encoding ensures that the system will be better guided along the various non-linearities of the trajectory.

Table III summarizes the steps of the complete procedure by which we iteratively test and re-estimate the system to improve and ensure local stability within the domain C.

IV. EXPERIMENTAL RESULTS

To validate the performance of the proposed method itself without blurring it with noise inherent to human demonstrations, we first tested its ability to reconstruct given theoretical dynamical systems. With a known system we may generate a clean training set, learn an approximation of the dynamics and further compare how well the learned dynamics approximate the real one.

Further, we verify the applicability of the method to robotics by teaching two robots manipulation tasks. We report on each of these next.

A. Learning Theoretical Dynamics

The method was validated to estimate four two-dimensional dynamical systems (*Systems 1-4*) and one three-dimensional dynamical system (*System 5*), each of them contains different number of attractors and exhibits different stability properties. In each case, we generated six trajectories using the theoretical dynamics and used these for training the GMM. When the dynamical system had more than one asymptotically stable attractor, trajectories were generated only in the subpart of the state space around one of them.

Note, the legend for Figures 7 - 10 is described in Figure 1. Each of the figures encompasses, in the first row, plots giving a general view of the original dynamics with vector fields (a) and three-dimensional phase plots (b-c), in the second row, a view of the GMM superimposed to the training data, and in the 3rd row, vector field (a) and phase plots (b-g) of the the estimated dynamics superimposed on the original dynamics.

System 1.

$$\begin{aligned} \dot{x}_1 &= -x_1 + 2x_1^2 x_2; \\ \dot{x}_2 &= -x_2. \end{aligned} \tag{13}$$

The system has a single locally asymptotically stable equilibrium point at the origin. We approximate the dynamics of this system in a region $[-4; 0] \times [0; 2]$, where it is locally asymptotically stable. Results are presented in the Figure 7.

System 2

$$\begin{aligned} \dot{x}_1 &= 700 - 2x_1 + 200 x_2 e^{\frac{25x_1 - 10^4}{x_1}}; \\ \dot{x}_2 &= 1 - x_2 - x_2 e^{\frac{25x_1 - 10^4}{x_1}}; \end{aligned} \tag{14}$$

The system has two equilibrium points – one asymptotically stable ($x_1 = 335$; $x_2 = 0.089$) and one unstable ($x_1 = 489$; $x_2 = 0.5$). We approximate the dynamics in the region $[0; 400] \times [-2; 2]$, where it is locally asymptotically stable. Results are presented in Figure 8.

System 3

$$\begin{aligned} \dot{x}_1 &= -x_2; \\ \dot{x}_2 &= x_1 - x_1^3 - x_2; \end{aligned} \tag{15}$$

The system has three equilibrium points - two unstable ($x_1 = -1; x_2 = 0$ and $x_1 = 1; x_2 = 0$) and one asymptotically

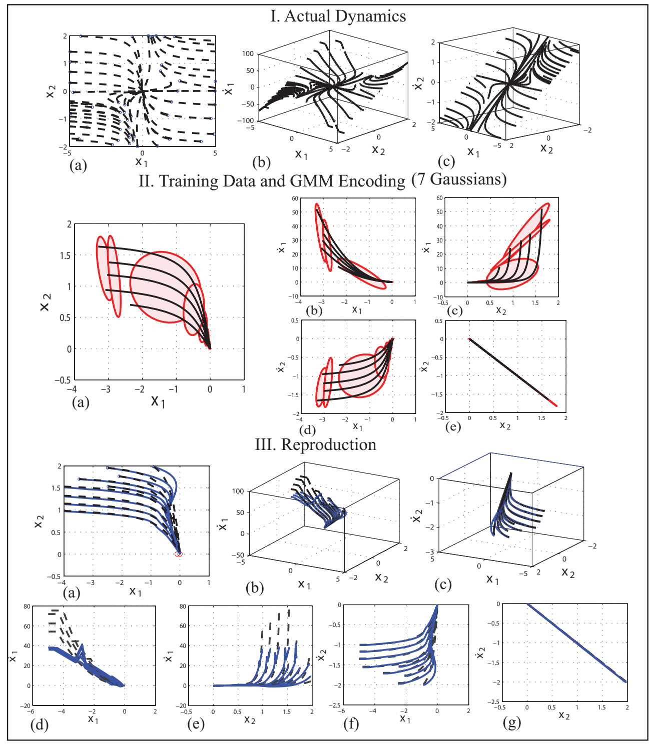

Fig. 7. *System 1*. The proposed method encodes this system with 7 Gaussians; the learned system exhibits good precision in the area covered by demonstrations, outside this area the precision is also admissible except for a region in the direct proximity to y-axis, where actual trajectories represent an excess curvature as approaching to the equilibrium, e.g., a trajectory starting at the bound $x_2 = 2$. In this region, a flat part of trajectories is reproduced well, though the steep parts that were not demonstrated are attracted towards the region covered by the training set.

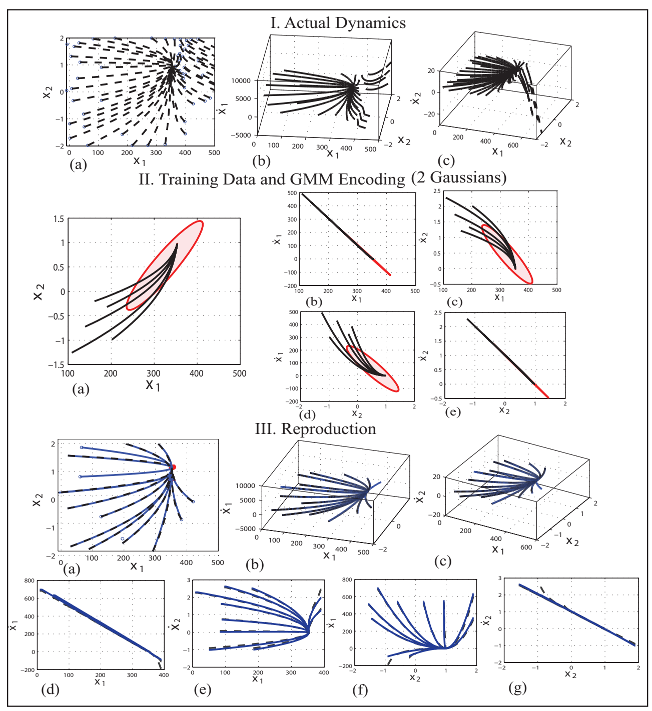

Fig. 8. *System 2*. As the behavior of the system in the considered area is relatively simple, 2 Gaussians are sufficient to achieve the good performance, even in areas unseen during demonstration. Interestingly, the learned dynamics is extrapolated very well beyond the area covered by the training set.

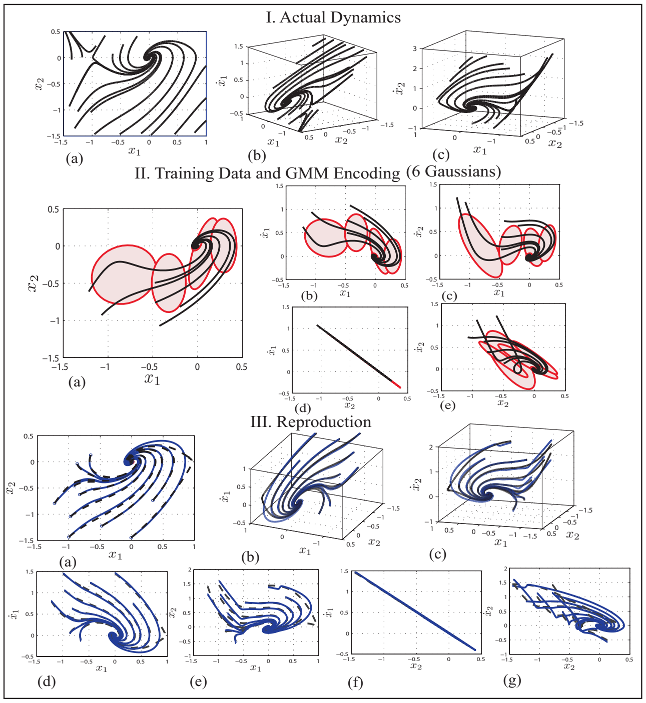

Fig. 9. *System 3*. Despite manifest non-linearity in the trajectories, the dynamics is successfully approximated with 6 Gaussians. Note, even unseen, circular shape trajectories (starting around $x_2 \approx 0$) are reproduced correctly in both position and velocities spaces.

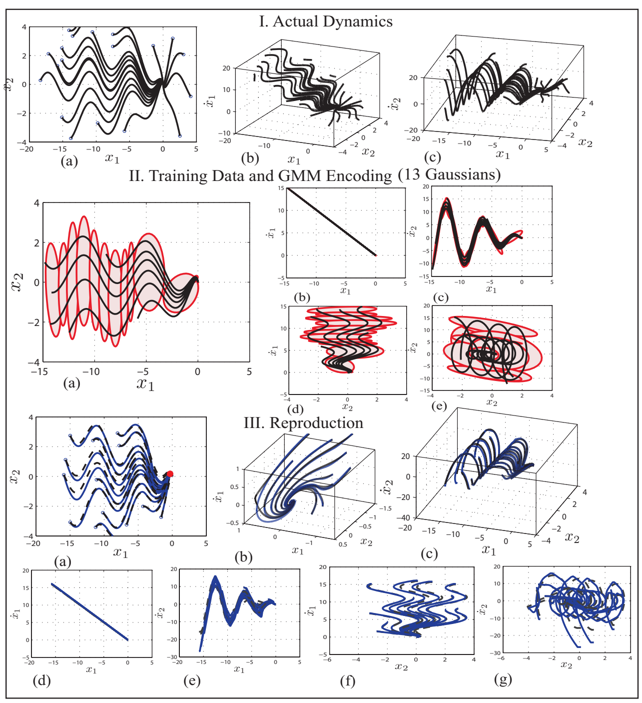

Fig. 10. *System 4*. The system is strongly non-linear, 13 Gaussians are necessary to achieve a good precision in the considered region. Complex dynamics and increased number of Gaussians lead to less strong generalization abilities of the method. Indeed, trajectories started beyond the region covered by the training set tend to depart from the real trajectories generated by the dynamics, it is particularly noticeable in the velocity space, see section III-(g). However, even in this non-trivial case the system generate admissibly good results from few demonstrations.

stable $x_1 = 0; x_2 = 0$. We approximate the dynamics of this system in a region $[-1.5; 1] \times [-1.5; 0.5]$, where it is locally asymptotically stable. Results are presented in Figure 9.

System 4

$$\begin{aligned} \dot{x}_1 &= -x_1; \\ \dot{x}_2 &= -x_1 \cos x_1 - x_2; \end{aligned} \tag{16}$$

The system exhibits strong nonlinearity due to the cosine term; the system is globally asymptotically stable and converges asymptotically to the origin. We approximate the dynamics of this system in a region $[-20; 0] \times [-4; 4]$. Results are presented in Figure 10.

System 5

$$\begin{aligned} \dot{x}_1 &= -x_1 - x_2 + x_3^2; \\ \dot{x}_2 &= x_1 + 10 \cos x_2 * x_2 - x_3^2; \\ \dot{x}_3 &= x_1 + 2x_2 - x_3; \end{aligned} \tag{17}$$

Locally asymptotically stable three-dimensional dynamics

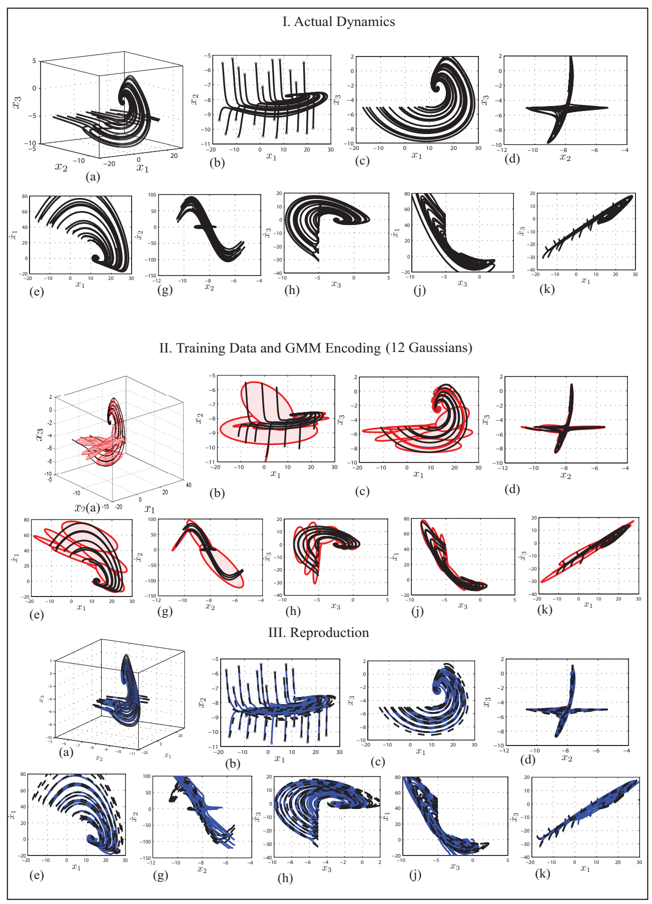

Fig. 11. *System 5*. Strongly non-linear 3D dynamics. In this case, a slight increase in a number of demonstrations allows for accurate approximation and generalization.

with a single attractor at $[12.98; -7.75; -2.5213]$. We approximate the dynamics of this system in a region $[-20; 30] \times [-11; -5] \times [-10; 2]$. Results of the learning process are presented in the Figure 11.

1) Quantification and discussion of results: Quantification of results achieved on both theoretical systems and actual robotic motions are presented in Table IV. As it can be seen all systems permit coarse representation with a relatively small number of Gaussians, moreover such a sparse representation achieves good precision in reproducing the actual dynamics. Furthermore, as shown in Figures 7-11, the system can generalize outside the training domain. This property is particularly useful for practical applications as this allows to predict the behavior of the system outside the region covered during training, hence reducing the amount of training data required. In the examples covered here, only 6 training trajectories were required in each case.

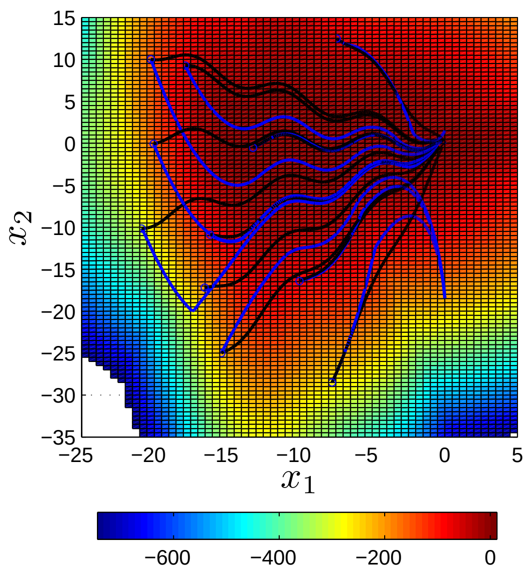

Fig. 12. Extrapolation properties of the GMMs encoding. A color map reflects changes in values of the the likelihood (18) of datapoints, the dark-red area represents an area of the most reliable inference regarding the velocity. For reconstructed trajectories starting outside this area, the deviation from the actual dynamics may be considerable. Interestingly, in the region of attraction of the origin, trajectories are strongly attracted towards a region covered by the training set. It is a useful property for practical applications as this allows to predict the behavior of the system outside the region covered during training, hence reducing the amount of training data required.

Note that, since the dynamics is learned from data covering only a subpart of the domain, it does not necessarily have the same attractor landscape and the region of attraction across the complete domain as the original system, even if it accurately approximates the original system locally. For example, in System 3, the original dynamical system has three equilibrium points, while its approximation has a unique asymptotically stable equilibrium. To overcome this, one may provide additional demonstrations covering dynamics in the neighborhood of the other equilibriums: Figure 15 presents results of learning the dynamics around the two different attractors of *System 5*. The demonstrations were provided in the neighborhood of the two asymptotically stable attractors; during learning, positions of two Gaussians were fixed on the attractors, and the algorithm was running to verify local asymptotical stability of both attractors. The regions of approximation C were analyzed separately for each attractors. The learned system managed to accurately grasp the complex dynamics, further, it allowed to separate the two flows of trajectories leading to different attractors based on the initial conditions of motion.

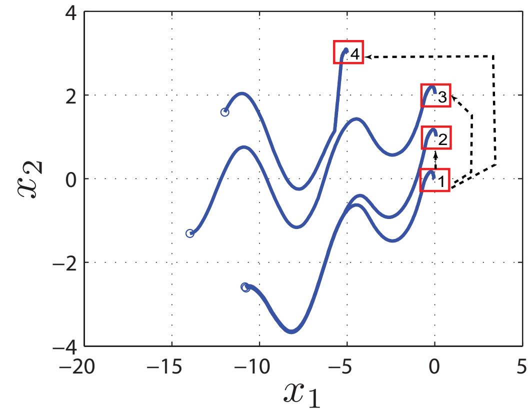

Fig. 13. Robustness to perturbations. The target is shifted several times (to positions 2, 3, 4) after the onset of motion.

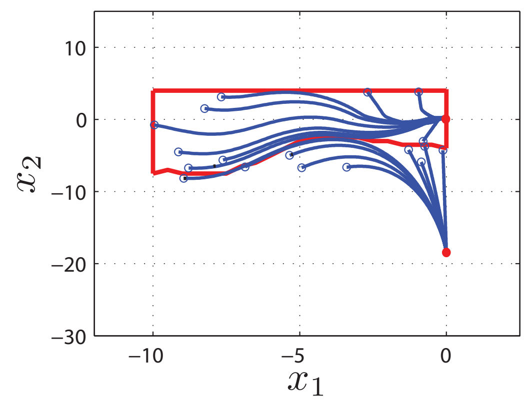

Fig. 14. Numerically estimated region of stability for the System 6; updated boundaries of the considered region taking into account boundaries of RA are plotted in red. An actual and spurious attractors are red circles. We extended a considered region from $[-20; 0] \times [-4; 4]$ to $[-10; 0] \times [-10; 4]$. Note, the numerical method estimated a lower bound that goes along a trajectory with a good precision. Other bounds were left unchanged, i.e., in the other directions the considered region does not cross boundaries of the region of attraction.

In addition to stability of reproduction, one should keep in mind that the considered region of applicability should not exceed a region where the likelihood of observing new data allows performing a confident inference regarding the velocity. In Figure 12 we depict how the likelihood changes beyond the region covered by the training set. Likelihood was computed as follows:

$$\mathbf{L}(\xi) = \log\left[\max_i h_i(\xi)\right]. \tag{18}$$

L gives a measure of the maximum probability of a point ξ to belong to any of the K Gaussians. The region where L exceeds a given threshold[8] represents the region where the system can still make a confident probabilistic inference. Note

[8]We took an empirically chosen threshold of -10.

TABLE IV
QUANTIFICATION OF RESULTS

System	Dimension	Number of demonstrations	Number of GMMs components	Number of iterations in model training[1]	Precision[2]
System 1	2	6	7	252	$8 \cdot 10^{-3}$
System 2	2	6	2	70	$3 \cdot 10^{-3}$
System 3	2	6	6	150	10^{-4}
System 4	2	6	13	475	10^{-1}
System 5	3	10	12	393	$8 \cdot 10^{-2}$
KATANA experiment	3	4	4	132	-[3]
HOAP experiment	3	4	5	160	-[3]

[1] The algorithm iterates until the change in the likelihood falls below 10^{-8}

[2] Precision is computed as a mean square error, on both seen and unseen trajectories, according to: $\frac{\sum_{i=1}^{M} \|\hat{\xi}^i - \xi^i\|^2}{M \cdot \Delta\xi}$, where $\hat{\xi}^i$ – learned trajectory ξ^i – theoretical trajectory, $\Delta\xi$ is an average amplitude of motion.

[3] Estimation of precision is non-applicable due to the presence of noise in the training data.

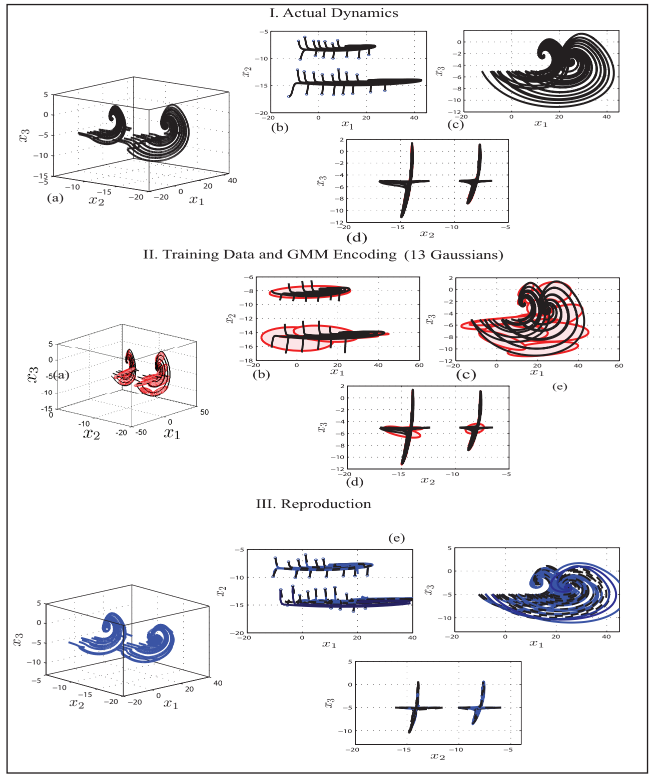

Fig. 15. Learning motion with two attractors. 3-dimensional trajectories are generated by *System 5* that displays a periodic behavior. Trajectories were demonstrated in the neighborhood of two asymptotically stable attractors. During the reproduction, the system managed to accurately reproduce dynamics around both attractors.

that all the trajectories that start in areas where L is too small will significantly depart from the real dynamics. This is due to the effect of the weights h_i associated to each Gaussian and how these influence the direction of the velocity vector: nearby the demonstrations, the influence of the closest Gaussian dominates that of all Gaussians, hence guiding closely the motion. However, far away from the demonstrations, the influence of all Gaussians becomes comparable and the resulting direction of velocity may point away from the signal.

As mentioned in the introduction, an inherent property of stable dynamical systems is their robustness to spatial and temporal perturbations. Figure 13 illustrates this aspect for one of the learned dynamical system, when the target is moved after the onset of the motion. As we see, the trajectories adapt smoothly to the change. Note, however, that the velocity profile may change abruptly when the perturbation occurs. To overcome this drawback it would be necessary to consider second-order dynamics.

As discussed previously, the GMMs encoding may result in spurious attractors outside the empirical stability domain C and in regions with low likelihood, see, e.g., Figure 12.

There are several reasons for the emergence of spurious attractors: first, the training set gives only a partial and noisy representation of the dynamics. Providing additional data in the regions around spurious attractors usually improves greatly performance. Second, the shape of the signal influence greatly stability. For instance, if the curvature of the trajectories changes smoothly, the spurious attractors, if any, will usually lie outside of the region of the confident inference, see Figure 12. However, if the system trajectories experience sharp changes in the curvature, as e.g., System 1, see the Figure 7, the likelihood of having spurious attractors in the considered region increases. By adding more Gaussians around the point with a sharp curvature one increases the guidance provided by the GMM and thus decreases the chances. By considering these practical shortcomings, one may improve a particular encoding to achieve the admissible performance.

V. APPLICATION TO ROBOT CONTROL

Further, we validate the method to learn the dynamics of motion of a robot endeffector when trained through human guidance. Here, the dynamics of motion becomes the control law that iteratively moves the robot's arm along a trajectory.

A. Encoding motion in the operational space

Since the framework we defined above does not make any assumption as to the type of variables to be used for training, we are unconstrained in our choice of variables for controlling a robot. Here, we choose to describe motions according to the following variables: the translation component of motion of the end-effector is described by a vector of Cartesian coordinates $\mathbf{x} \in \mathbb{R}^3$.

Each demonstrated trajectory is, thus, represented by the following dataset: $D = \{\mathbf{x}_t, \dot{\mathbf{x}}_t\}_{t=1}^{M}$, where M is the number of datapoints in a trajectory. To reproduce a task, we first learn an estimate of the dynamical system using the method described

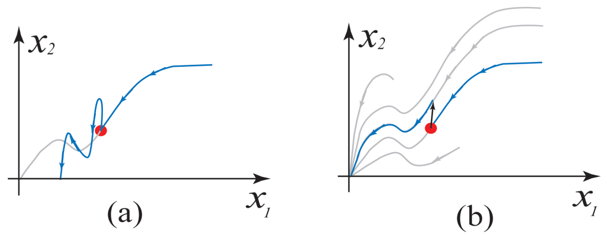

Fig. 16. (a) If a trajectory in the operation space passes through non-reachable joint positions IK may return velocity in the operation space that sends a robot too far from original trajectory, so linear assumptions of approximation of kinematics does not satisfy and overall trajectory tracking will fail. (b) In the case of motion encoding with a dynamical system, after perturbation the robot will not try to return to the previous trajectory violating the linear approximation of kinematics, instead the dynamical system will generate other trajectory from the point where the robot occurs.

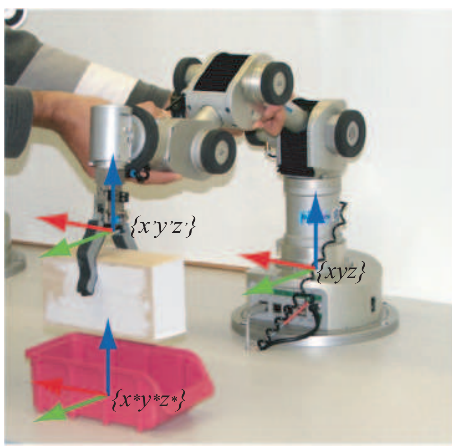

Fig. 17. We encode tasks in a referential located at the target and moving with it$\{x^*y^*z^*\}$; this referential is expressed in the fixed global referential $\{xyz\}$(usually we choose one attached to static parts of a robot). Actually, the motion of the robot end-effector is expressed as moving a referential associated with the end-effector $\{x'y'z'\}$.

in Section III-A and then use the Moore-Penrose pseudo-inverse to compute the corresponding joint angles. Table V summarizes the steps of the reproduction algorithm.

B. Set-up

We validated the above method in two practical tasks, see Figure 18 where a human teacher guides the robot through the motion. We also implemented the learned 3-dimensional *System* 5, as a control law for our robot. To demonstrate the generic character of the approach we ran experiments with two different robotic platforms: a 6 degrees of freedom industrial-like KATANA arm from Neuronics and a 4 degrees of freedom robot arm of the humanoid robot HOAP-3 from Fujitsu.

C. Experiments with KATANA

The first experiment consists in the KATANA putting an object into a container. Here, the KATANA arm was taught to put a rectangular wooden brick into a rectangular container; see fig.18-left.

TABLE V
ON-LINE TASK REPRODUCTION

1	Assume that a controller $\hat{f}_{\mathbf{x}}$ has been learned, the robot is thus ready to reproduce a task
2	Detect a target position in the global referential $\{xyz\}$; see Figure 17: $\mathbf{x}^*$
3	Recompute the current position of an end-effector in the target referential $\{x^*y^*z^*\}$: $\mathbf{x}_0$
4	**LOOP** until the target position is reached
6	infer the velocity for the next iteration t through GMR Eq.9: $\hat{\dot{\mathbf{x}}}_t$ $\hat{\dot{\mathbf{x}}}_t = \sum_{k=1}^{K} h_{k,\mathbf{x}}(\mu_{k,\dot{\mathbf{x}}} + \Sigma_{k,\dot{\mathbf{x}}\mathbf{x}}\Sigma_{k,\mathbf{x}}^{-1}(\mathbf{x} - \mu_{k,\mathbf{x}}))$
8	Solve the Inverse Kinematics problem to find: $\dot{\mathbf{x}}_t, \dot{\boldsymbol{\theta}}_t$
9	compute a new position $\mathbf{x}_t, \boldsymbol{\theta}_t$
10	**END**

In the second experiment, the KATANA was controlled with *System 5* with the origin of the system positioned on an arbitrary object. This experiment meant to test the ability of the learned system to generalize to context unseen during training and to quickly adapt to perturbations.

D. Experiments with HOAP-3

The clench of the HOAP-3 is rather small, therefore it can grasp only thin objects. In this task the robot had to grasp a box which is thin along one dimension, so the robot should follow a specific path to properly position its hand; see fig.18-right.

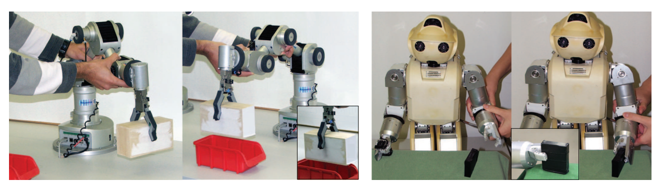

Fig. 18. Set-up of the experiments. Left: KATANA puts a wooden brick into the container, to achieve the task the robot should lift the brick and move it following an elevating trajectory. Right: HOAP-3 grasps a box, to achieve the task HOAP should approach the box with a specific orientation and than lower its arm, as the clench is small, see small figure in the corner.

During training, the robots were shown the tasks 5 times by a human user guiding their arms. Values of the robots joints were recorded during this passive motion and used for reconstructing the position of the end-effector.

E. Results of Learning Dynamics from Motion Data

After training, we tested the system by requesting the robots to reproduce the tasks in various conditions. The results of the experiments are summarized in Figures 19-22.

To test the generalization abilities and the stability to perturbations we performed experiments in different conditions, by changing the starting positions of the robots and shifting the container (for the KATANA's experiment) or the box (for the HOAP-3's experiment). Results are presented in Figure 22; in both experiments learning of position control was successful and the robots all reached successfully the targets and accomplished the tasks.

Results of generalization for the second experiment with KATANA reproducing *System 5* are presented in Figure 20 - II. The area where demonstrations were provided is depicted in Figure 20 - II (b) with red squares. The system further allowed to reproduced the motion starting from any position of the sub-space of the workspace, depicted in grey. Note, that even few demonstrations provide good generalization properties.

The ability to generate a trajectory from arbitrary initial position to the target with a relevant velocity profile is a strong point of encoding motion with Dynamical Systems in the state-space, furthermore it provides real-time adaptation to perturbations in the position of the target. The Figure 20-I presents results of tracking a marked object mapped into the attractor of the dynamical system. After shifts of the target, the robot finally reaches the object following the demonstrated position and velocity profile.

VI. Discussion and Future Work

Below, we discuss the major hypotheses postulated in this work together with possible alternative solutions.

A. *Multi-dimensional systems, first order dynamics*

The method proposed here allows learning of non-linear multivariate dynamics where the correlation between the variables is important. Other works on dynamical control consider each degree of freedom separately, hence discarding information pertaining to correlation across the joints. While storing correlations across the joints is costly (in GMM, it forces one to compute the complete covariance matrix, rather than computing only the diagonal elements), it is advantageous as correlations contain features characteristic of the motion. For instance, in bimanual coordination tasks in which left and right arms should follow different dynamics while doing so in coordination [Gribovskaya and Billard, 2008], embedding the correlations in the representation ensures the reproduction of both the dynamics of each arm and the correlations across the arms. Furthermore, learning correlation between a multivariate signal and its derivatives allows to considerably decrease a number of Gaussians required to accurately encode the training dataset.

While we started with the hypothesis that the control law followed a first order dynamics, the method proposed here may be extended to learn higher-order dynamics (as higher-order systems can always be expressed in the canonical form as a set of first-order systems). That is particularly relevant for applications where it is necessary to control the acceleration profile. We intend to address this problem in future work.

Potential difficulties concerning shifting into higher-order derivatives that can be envisioned, are associated with the increased dimensionality of a resultant statistical problem. With an increase in the number of dimensions, a stable approximation would require more training data or need to introduce certain heuristics to partially decouple the problem into a set of systems with lower dimensions.

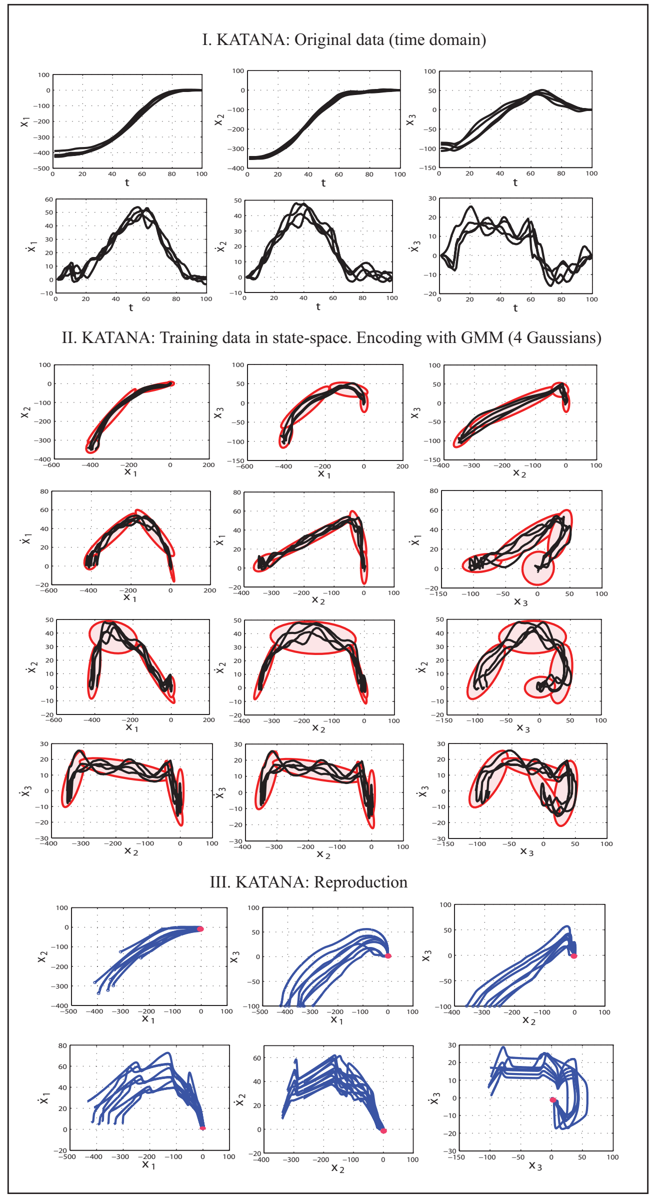

Fig. 19. KATANA experiment 1: Results of encoding and reproduction of the experiment where KATANA had to put a brick into a container.

B. *Time independency* vs *time dependency*

In this paper, we advocate that time-independent encoding in the state-space offers more robust representation in comparison to traditional time-dependent encoding. Results confirmed that for a certain range of motions, the state-space representation is indeed highly robust to spatial and temporal perturbations. Moreover, it allows to reproduce tasks even in unseen parts of the workspace.

Yet, certain motions, such as those requiring the synchronization with an external dynamics, should be encoded using a time-dependent representation or, if the external dynamics is known, using an explicit parametrical coupling of two time-independent dynamics, such as that done in [Ijspeert et al., 2001]. Another limitation of the time-independent representation relates to the possibility of encoding compound motions: in this case, the whole motion may be segmented into a set

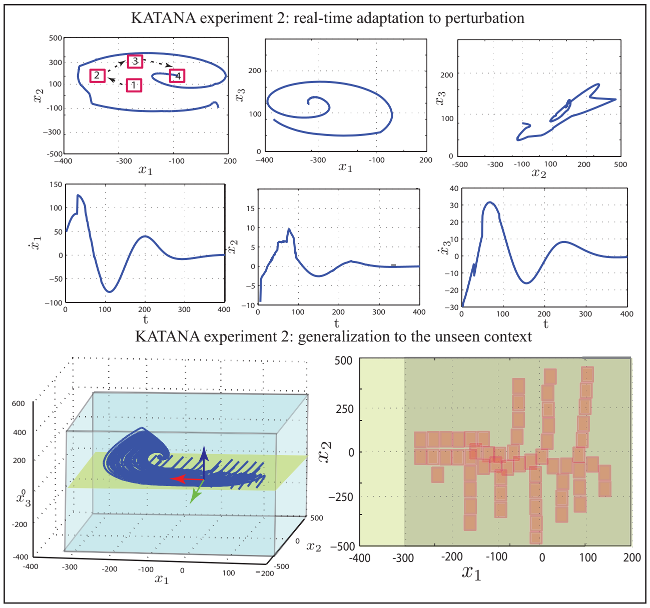

Fig. 20. KATANA experiment 2: *I. Real-time adaptation to perturbations*. The target was consequently shifted from the position 1 to the position 4. First row: trajectory of the robot's end-effector; second row: velocity profile. *II. Generalization to the unseen context*. (a) The workspace of KATANA is highlighted by the blue box, the reproduction was systematically tested starting robot from positions on the yellow plane. (b) The starting plane from the robot workspace is in yellow. The robot was required to reproduce the motion from points monotonically covering the part of the starting plane (in grey). For comparison, the part of space, where the demonstrations were provided is in pink. Notice, that demonstrations are sparse, but the system manages to generalize to other parts of the workspace.

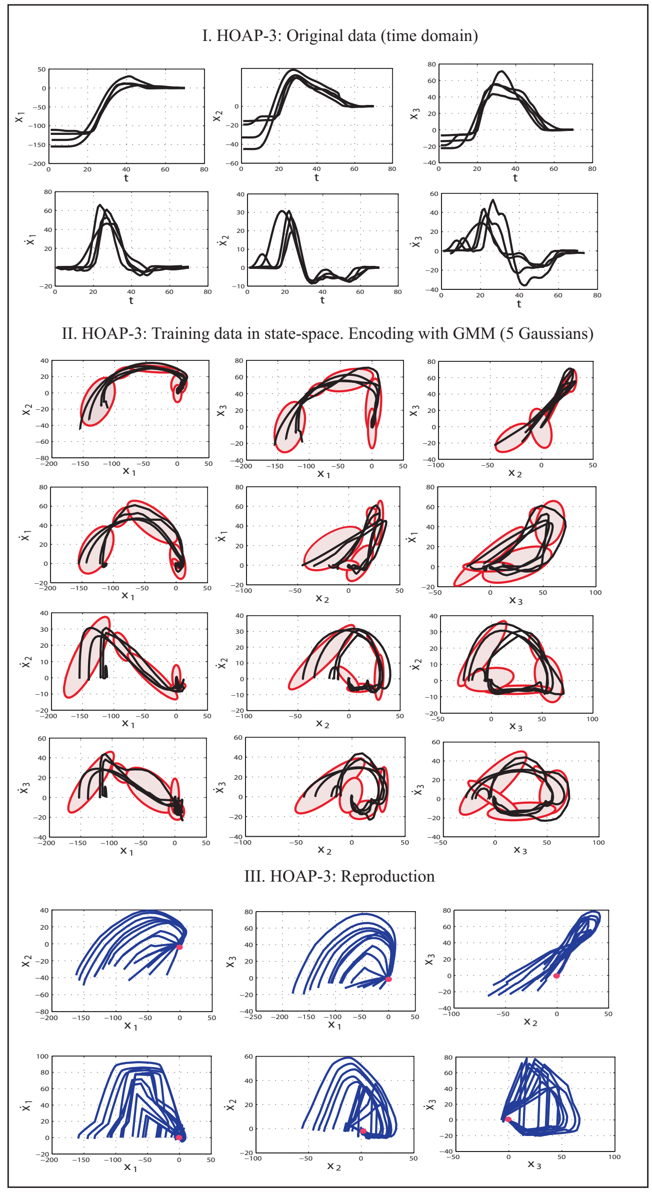

Fig. 21. HOAP-3 experiment: Results of encoding and reproduction of the experiment where HOAP-3 had to grasp a box.

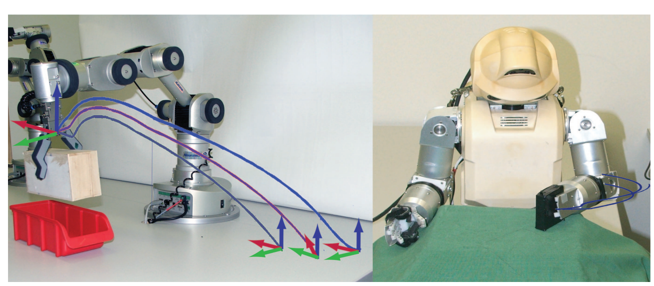

Fig. 22. The results of reproduction of dynamically generated trajectories on the robots. To check the generalization abilities of the learned dynamics the trajectories were reproduced from different initial positions.

of simpler ones governed by a single attractor. However, the problem of how to transit across these systems remains an open issue.

C. Kinematic controller

In the experiments reported here, control of the robot was purely kinematical, encoding the desired kinematic trajectories, but not taking into consideration the dynamical properties (actual torques) of the robot limbs. An additional control step was then necessary to convert positions into motor commands by means of the inverse dynamics (KATANA) or a PID controller (HOAP-3). Learning the inverse dynamics, while a highly value topic in itself, is beyond the scope of the present paper. Further, considering that many of the current robotic platforms are position-controlled, while providing position feedback in real-time, the proposed approach is thus valid for a large set of applications.

D. Choice of statistical framework

GMMs being a global statistical techniques (by opposition to local non-parametric methods such as LWPR, GPR) was shown to be suitable for estimating dynamics from sparse demonstrations, that are typical of programming by demonstration applications. However, neither GMMs nor LWPR and GPR ensure stability of a learned approximation. Here, we proposed an algorithm that leads to local asymptotical stability and gradually improves the quality of the approximation while widening the region of applicability C. Potentially, the same procedure may be adopted for other statistical frameworks. However, the accuracy of the approximation may significantly vary depending on a particular choice.

One should note that EM is more computationally expensive than LWPR with a number of iteration steps during training of $O(K \cdot M \cdot N)$ in comparison to $O(N)$. Both of these however remain small in comparison to GPR. Similarly to LWPR and in contrast to the GPR-based methods, GMR's computational costs for the retrieval procedure are low and increase linearly with the number of parameters. Additionally, GMM-based models result in much less parameters due to the coarse representation.

E. Real-time adaptation to perturbation vs traditional planners

One of the strengths of the proposed approach is its ability to cope with perturbations in real-time. By perturbation we referred to unexpected changes in the positions of the attractor or of the robot's joints during motion. We demonstrated how the learned dynamics with a position of an object mapped into an attractor can successfully track the object. Such a flexibility combined with the guarantee of ultimately reaching the object is one of the major advantages of the proposed method in comparison with traditional planners [Yokoi et al., 2009, Yoshida et al., 2008, Diankov and James Kuffner, 2007, Kuffner et al., 2002]. One should emphasize that planners, in turn, are advantageous when the environment is known and for providing mechanisms for obstacle avoidance. The latter is, however, achieved by introducing a heuristic-based cost function that penalizes certain directions. Potentially, our approach may be combined with such a cost function that perturb an output of a learned dynamical system pushing it away from obstacles.

Note, that our system though introduces certain hypotheses, still remains rather generic regarding tasks it may reproduce, furthermore, it may work with limited and inaccurate information about the environment, as it does not require any costly replanning. At the same time, to benefit from optimal planning and capacity for obstacle avoidance, one should provide an algorithm with precise information regarding objects in the workspace and introduce certain task-related heuristics to improve convergence.

F. Single vs *several attractors*

A further hypothesis pertaining to the work presented here was the idea that the dynamical system to be discovered had a single or several known fixed point attractors. This can be considered as a limitation, as a dynamics may be governed by the existence of more complex orbits than merely fixed points. For example, an arbitrary free motion may have a particular curve in space as attractor. The applicability of the proposed method in this case will mostly depend on the quality of training data; further no stability can be guarantee. Procedures for ensuring stability of complex orbits may substantially widen the class of motion under consideration, covering dancing or sport motions that are usually characterized by the existence of certain curves to which all trajectories converge.

G. Training data

The generalization properties of dynamical controllers directly depend on the quality of training data; the aspect common to all statistical learning methods. It might be compensated in different ways: 1) by providing an exhaustive set of accurate demonstrations; 2) by allowing a robot to explore on its own (considered in Reinforcement Learning [Guenter et al., 2007]); 3) by providing more variability in a limited set of demonstrations (the problem has been discussed in [Calinon and Billard, 2007]). The first option does not agree with a requirement of user-friendliness of teaching interfaces, as a number of demonstrations should be kept bearable for a user; the second approach may require additional time; therefore, we concentrate on improving quality of demonstrations by introducing more variability into a small set of demonstrations.

H. Kinesthetic teaching

For demonstrating tasks we used the kinesthetical teaching approach that consists of directly demonstrating the task using a robot's own body. One of advantages of this approach is that the human can feel limitations of the robot's architecture and adapt his/her intuition about an optimal or efficient motion accordingly. Although we actively exploit this learning paradigm, other approaches such as vision-based learning are also widely used and can be more intuitive for humans. Our system may be applied to the motion data obtained through different modalities.

I. Practical consideration

From a practical point of view, mapping position of manipulated objects into attractors of Dynamical Systems considerably improves the precision of motion at a target and therefore allows considering prehensile tasks in the framework of Programming by Demonstration; where so far generation of large-scale motions has been addressed.

The approach was shown to be generic in that it did not depend on the particular geometry of the robot's arm, nor on the particular variables to be learned. Indeed, it could be successfully implemented to control robot arms with different geometries and for learning the dynamics of different variables inherent to position and orientation control. **Source code and supplementary material is available at http://lasa.epfl.ch/elena/learning-dynamics.htm**

VII. Summary

In this paper, we proposed a method for learning a nonlinear multi-dimensional dynamics of motion through statistically encoding demonstrated data with Gaussian Mixtures. Further, we addressed the problem of ensuring stability of a resultant control law: first, we formulated conditions that parameters of GMMs should satisfy to guarantee local asymptotical stability of an attractor, then we proposed a numerical procedure to verify boundaries of the region of applicability where the control law can be securely applied.

To test the method, we conducted two types of experiments: 1) learning theoretical dynamics with known mathematical forms to estimate the accuracy of approximation and 2) learning dynamics of manipulation tasks recorded with two different robotic platforms to assess the applicability of the approach to the noisy data. In all experiments the system demonstrated good results in terms of high accuracy during reproduction, ability to generalize motions to unseen contexts, and ability to adapt on-the-fly to spatio-temporal perturbations. We also showed how the system can encode more than one attractor, successfully reproducing each separate dynamics locally around each attractor and separating the flows leading to the different attractors.

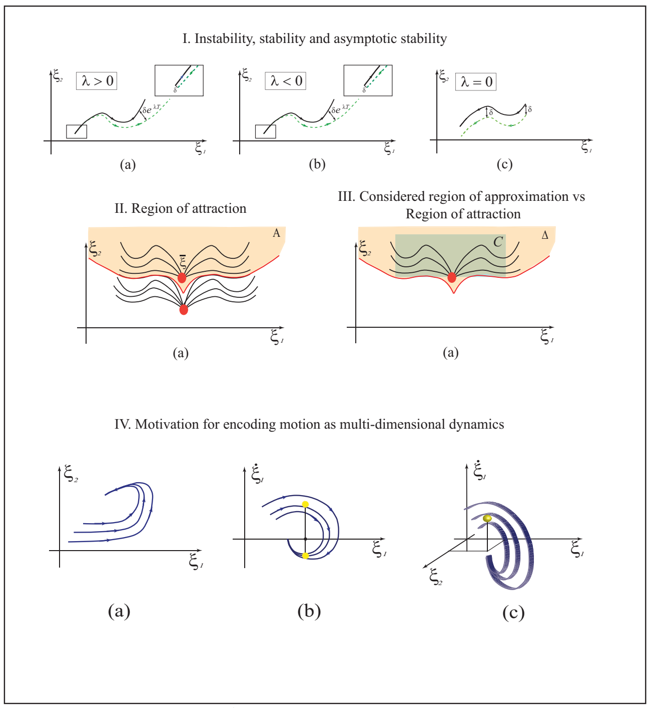

Fig. 23. Appendix II. Geometrical illustration of stability and multi-dimensional correlation in the state-space. I. ***Stability problem***: stability of a dynamical system is defined by a maximum value of its *Lyapunov exponent* λ (in the linear case, it coincides with eigenvalues of a control matrix). (a) In systems with *negative* Lyapunov exponents volume between trajectories contracts; (b) In systems with *positive* Lyapunov exponents two arbitrary near trajectories diverge from each other exponentially fast. In the linear case, one may easily find Lyapunov exponents and estimate the global behavior of the overall system. In the non-linear case, the system may have different Lyapunov exponents in different parts of the state-space, moreover, non-linearities make analytical investigation of properties particularly tedious. IV. ***Multi-dimensional dynamics*** Analyzing dynamics of vector-valued timeseries requires their encoding in multi-dimensional state-spaces. Generally, one cannot unambiguously decouple dynamics of each dimension. Consider a simple 2D motion in Figure II-(a), the phase-space of this motion in $\{\dot{x}_1, x_1\}$ is in Figure II-(b): for each value x_1 there exist two different values of velocity, therefore, it is not possible to unambiguously encode dynamics of motion as two decoupled system $\dot{x}_1 = f_1(x_1)$, $\dot{x}_2 = f_2(x_2)$. However, if one look at the dependence $\dot{x}_1 = f(x_1, x_2)$ depicted at Figure II-(c) this ambiguity can be easily eliminated. This problem is know in the literature on Dynamical Systems as a problem of searching for a *minimum embedding dimension*. In this particular example, the minimum embedding dimension is 4 (x_1, $\dot{x}_1$, x_2, $\dot{x}_2$). Alternatively, one may argue that in this case we may avoid an ambiguity and separate dimensions encoding $\ddot{x}_1 = f_1(x_1, \dot{x}_1)$, though it is possible in this particular case, it will lead to the necessity to analyze 5 state variables (x_1, $\dot{x}_1$, $\ddot{x}_1$, x_2, $\dot{x}_2$). Furthermore, to preserve a spatial correlation pattern between x_1 and x_2 the decoupled systems should be synchronized by an external mechanism.

TABLE VI
APPENDIX I. COMPARISON OF THE PROPOSED METHOD WITH [IJSPEERT ET AL., 2001]

GMR-based method proposed in this paper:

a single multidimensional system is running to control several DOFs

$$\dot{\mathbf{x}} = \hat{f}(\mathbf{x})$$
$$\hat{f}(\mathbf{x}) \triangleq \sum_{k=1}^{K} h_k(\mathbf{x})(\mu_{k,\dot{\mathbf{x}}} + \Sigma_{k,\dot{\mathbf{x}}\mathbf{x}}\Sigma_{k,\mathbf{x}}^{-1}(\mathbf{x} - \mu_{k,\mathbf{x}}))$$

where $\mathbf{x} \in \mathbb{R}^N$; $\Sigma_{k,\dot{\mathbf{x}}\mathbf{x}},\ \Sigma_{k,\mathbf{x}} \in \mathbb{R}^{N\times N}$ are estimated matrices
$\mu_{k,\dot{\mathbf{x}}},\ \mu_{k,\mathbf{x}} \in \mathbb{R}^N$ are estimated vectors

---

LWPR-based method proposed in [Ijspeert et al., 2001] (see also Figure 2):

the velocity along each DOF $\dot{x}$ is defined by a linear combination of two velocities $\dot{z}$ and $\dot{\nu}$, according to:

$$\dot{x} = \dot{z} + \hat{f}^*(\nu)\dot{\nu}$$

$$\hat{f}^*(\nu) \triangleq \frac{\sum_{k=1}^{K} \Psi_k(\nu)\omega_k}{\sum_{k=1}^{K} \Psi_k(\nu)}$$

where $x,\ z,\ \nu \in \mathbb{R}$
$\Psi_k(\nu) = \exp\frac{(\nu - c_k)^2}{2\sigma_k^2}$, $\omega_k \in \mathbb{R}$.

The variables z and ν are governed by two dynamical system:

(S1)	(S2)
$\ddot{\nu} = \alpha_v(\beta_v(g - \nu) - \dot{\nu})$	$\ddot{z} = \alpha_z(\beta_z(g - y) - \dot{z})$ $\dot{y} = \dot{z} + \hat{f}^*(\nu)\dot{\nu}$
where $g,\ v \in \mathbb{R}$, $\alpha_v,\ \beta_v \in \mathbb{R}$ are known constants	where $y, z \in \mathbb{R}$; $\alpha_z, \beta_z \in \mathbb{R}$ are known constants

---

Comparison between GMR-based $f(\mathbf{x})$ and LWPR-based $f^*(\nu)$:

Function $f^*(\nu)$ represents a uni-dimensional simplified version of a function $f(\mathbf{x})$, indeed, weights $h_k(\mathbf{x})$ have the same form of Gaussians as $\Psi_k(\nu)$, further instead of introducing a variable components $(\mu_{k,\dot{\mathbf{x}}} + \Sigma_{k,\dot{\mathbf{x}}\mathbf{x}}\Sigma_{k,\mathbf{x}}^{-1}(\mathbf{x} - \mu_{k,\mathbf{x}}))$, LWPR considers merely constants ω_k, which is equivalent to using solely $\mu_{k,\dot{\mathbf{x}}}$.

Weights ω_k are tuned so to minimize a mean-square error between the velocity $\dot{y}$ and a demonstrated velocity profile.

Note, that according to the LWPR-based method [Ijspeert et al., 2001] the function $\hat{f}^*(\nu)$ modulating the velocity $\dot{x}$ does not depend on the actual position x, but instead depends on the internal state ν and, therefore, does not introduce a feedback loop. Practically it means that the only term adapting during perturbations is $\dot{z}$, while $\hat{f}^*(\nu)\dot{\nu}$ remains the same and may deform a trajectory.

The system (S1) is a spring and damper system, which attracts a trajectory ν towards the target g following a straight line path.

The system (S2) is a perturbed spring and damper system: initially it starts to go exactly as the system (S1), but due to the perturbed velocity $\dot{y}$, it departs from the straight line trajectory ν; the source of perturbation is a component $\hat{f}^*(\nu)\dot{\nu}$.

VIII. ACKNOWLEDGMENTS

This work was supported by the European Commission through the EU Projects FEELIX-GROWING (FP6-IST-045169) and ROBOT@CWE (FP6-034002).

The authors warmly thank Prof. Auke Jan Ijspeert for advices and support toward the research work developed here.

REFERENCES

J. Aleotti, S. Caselli, and M. Reggiani. Evaluation of virtual fixtures for a robot programming by demonstration interface. *Transactions on Systems, Man, and Cybernetics*, 35(4): 536–545, 2005.

R. Andersson. Aggressive trajectory generator for a robot ping-pong player. volume 9, pages 15–21, Feb 1989. doi: 10.1109/37.16766.

M. Aoki. *State Space Modeling of Time Series*. Springer-Verlag, 1990.

X. Bai, X.-S. Yang, and H. Li. Estimates of the region of attraction of continuous-time cascade systems. *Journal of Mathematical Control and Information*, 24(4)::483–491, 2007.

R. Bellman. *Dynamic Programming*. NJ:Princeton Univ. Press, 1957.

A. Billard, S. Calinon, R. Dillmann, and S. Schaal. Robot Programming by Demonstration. In *Handbook of Robotics*, volume chapter 59. MIT Press, 2008.

M. D. Buhmann. *Radial Basis Functions: Theory and Implementations*. Cambridge University Press, 2003.

S. Calinon and A. Billard. What is the Teacher's Role in Robot Programming by Demonstration? - toward Benchmarks for Improved Learning. *Interaction Studies. Special Issue on Psychological Benchmarks in Human-Robot Interaction*, 8 (3), 2007. doi: NA.

S. Calinon and A. Billard. A Probabilistic Programming by Demonstration Framework Handling Constraints in Joint Space and Task Space. In *In the Proceedings of the International Conference of Intelligent Robots and Systems*, 2008. URL http://iros2008.inria.fr/.

S. Calinon, F. Guenter, and A. Billard. On Learning, Representing and Generalizing a Task in a Humanoid Robot. *IEEE transactions on systems, man and cybernetics, Part B. Special issue on robot learning by observation, demonstration and imitation*, 37(2):286–298, 2007. doi: NA.

T. L. Carroll. A nonlinear dynamics method for signal identification. *Chaos*, 17, 2007.

F. Chamroukhi, A. Same, G. Govaert, and P. Aknin. Time series modelling by a regression approach based on a latent process. *Neural Networks.*, 22:593–602, 2009.

M. B. Crutchfield J.P. Equation of motion from a data series. *Complex Systems*, 1:417–452, 1987.

K. D., W. Takano, and Y. Nakamura. Incremental learning, clustering and hierarchy formation of whole body motion patterns using adaptive hidden markov chains. *The International Journal of Robotics Research*, 27(7):761–784, 2008.

M. Deisenroth, C. Rasmussen, and J. Peters. Gaussian process dynamic programming. *Neurocomputing*, 72:1508–1524, 2009.

Y. Demiris and B.Khadhouri. Hierarchical attentive multiple models for execution and recognition of actions. *Robotics and Autonomous Systems*, 54:361–369, 2006.

A. Dempster, N. Laird, and D. Rubin. Maximum likelihood from incomplete data via the em algorithm. *Journal of Royal Statistic Society*, 39:1–38, 1977. doi: http://dx.doi.org/10.2307/2984875. URL http://dx.doi.org/10.2307/2984875.

R. Diankov and J. James Kuffner. Randomized statistical path planning. In *Proceedings of IEEE/RSJ International Conference on Robots and Systems (IROS)*, 2007.

K. Dixon and P. Khosla. Trajectory representation using sequenced linear dynamical systems. *Robotics and Automation, 2004. Proceedings. ICRA '04. 2004 IEEE International Conference on*, 4:3925–3930, 26-May 1, 2004. ISSN 1050-4729.

R. Genesio, M. Tartaglia, and A. Vicino. On the estimation of asympototic stability regions: state of the art and new proposals. *IEEE Transaction on Automatic Control*, 30(8): 1985, 1985.

P. Giesl. Construction of a local and global lyapunov function using radial basis functions. *Journal of Applied Mathematics*, 73(5)::782–802, 2008.

E. Gribovskaya and A. Billard. Combining Dynamical Systems Control and Programming by Demonstration for Teaching Discrete Bimanual Coordination Tasks to a Humanoid Robot. In *IEEE/ASM International Conference on Human-Robot Interaction*, 2008.

F. Guenter, M. Hersch, S. Calinon, and A. Billard. Reinforcement Learning for Imitating Constrained Reaching Movements. volume 21, pages 1521–1544, 2007. doi: NA.

W. Hardle. *Smoothing Techniques with Implementation in Statistics*. NY:Springer, 1991.

M. Hersch, F. Guenter, S. Calinon, and A. Billard. Dynamical System Modulation for Robot Learning via Kinesthetic Demonstrations. *IEEE Transactions on Robotics*, 1, 2008. Accepted.

P. A. Hinrichsen D. *Mathematical Systems Theory*. Springer Berlin, 2000.

J. Hwang, R. Arkin, and D.-S. Kwon. Mobile robots at your fingertip: Bezier curve on-line trajectory generation for supervisory control. *Intelligent Robots and Systems, 2003. (IROS 2003). Proceedings. 2003 IEEE/RSJ International Conference on*, 2:1444–1449, Oct. 2003.

A. Ijspeert and A. Crespi. Online trajectory generation in an amphibious snake robot using a lamprey-like central pattern generator model. *Robotics and Automation, 2007 IEEE International Conference on*, 1:262–268, April 2007. ISSN 1050-4729. doi: 10.1109/ROBOT.2007.363797.

A. Ijspeert, J. Nakanishi, and S. Schaal. Trajectory formation for imitation with nonlinear dynamical systems. *Intelligent Robots and Systems, 2001. Proceedings. 2001 IEEE/RSJ International Conference on*, 2:752–757 vol.2, 2001.

A. Ijspeert, J. Nakanishi, and S. Schaal. Learning rhythmic movements by demonstration using nonlinear oscillators. In *In Proceedings of the IEEE/RSJ Int. Conference on Intelligent Robots and Systems (IROS2002*, pages 958–963, 2002.

H. Khalil. *Nonlinear systems*. Prentice Hall Upper Saddle River, NJ, 1996.

J. Kuffner, S. Kagami, and I. M. I. H. Nishiwaki, K. Dynamically-stable motion planning for humanoid robots. *Autonomous Robots*, 12(1):105–118, 2002.

E. T. Y. Lee. Comments on some b-spline algorithms. *Computing (Springer-Verlag)*, 36:229–238, 1986.

L. Ljung. State of the art in linear system identification: Time and frequency domain methods. *Proceeding of the 2004 American Control Conference*, 1:650–661, 2004.

M. Loccufier and E. Noldus. A new trajectory reversing method for estimating stability regions of autonomous non-linear systems. *Nonlinear Dynamics*, 21(3):265–288, 2000.

G. McLahlan and D. Peel. *Finite Mixture Models*. NY:Wiley, 2000.

A. Moore. *Efficient memory-based learning for robot control*. PhD thesis, University of Cambridge, 1990.

H.-G. Muller. *Nonparametric Regression Analysis of Longitudinal Data*. Berlin:Springer, 1988.

D. Nguyen-Tuong, M. Seeger, and J. Peters. Local gaussian process regression for real time online model learning and control. In *Proceedings of the Conference on Advances in Neural Information Processing Systems*, 2008.

J. Park and I. Sandberg. Universal approximation using radial-basis-function networks. *Neural Computation*, 3(2):246 – 257, 1991.

M. Priestley. State-dependent models: A general approach to non-linear time series analysis. *Journal of Time Series Analysis*, 1, 1980.

L. Righetti, J. Buchli, and A. Ijspeert. Dynamic hebbian learning in adaptive frequency oscillators. *Physica D*, 216(2):269–281, 2006. URL http://dx.doi.org/10.1016/j.physd.2006.02.009.

R. W. L. Ryoung K. Lim, Minh Q. Phan. State-space system identification with identified hankel matrix. Technical report, Department of Mechanical and Aerospace Engineering Technical Report No.3045, Princeton University, Princeton, NJ., 1998.

S. Schaal and C. Atkeson. Robot juggling: implementation of memory-based learning. *Control Systems Magazine, IEEE*, 14(1):57–71, Feb 1994. ISSN 0272-1708. doi: 10.1109/37.257895.

S. Schaal and C. G. Atkeson. Constructive incremental learning from only local information. *Neural Computation*, 10(8):2047–2084, 1998. URL citeseer.ist.psu.edu/schaal97constructive.html.

S. Schaal, S. Kotosaka, and D. Sternard. Nonlinear dynamical systems as movement primitives. In *Proceedings of the International Conference on Humanoid Robotics*, 2001.

S. Schaal, A. Ijspeert, and A. Billard. Computational Approaches to Motor Learning by Imitation. *Philosophical transactions: biological sciences*, 358(1431):537–547, 2003. doi: NA.

G. Schoner and C. Santos. Control of movement time and sequential action through attractor dynamics: A simulation study demonstration object perception and coordination. In *Symposium on Intelligent Robotic Systems*, 2001.

D. Sternad and D. Schaal. Segmentation of endpoint trajectories does not imply segmented control. *Experimental Brain Research*, 124:118–136, 1999.

H. G. Sung. *Gaussian Mixture Regression*. PhD thesis, Rice Univercity, Huston, Texas, 2004.

H. Tomohisa, W. M. Haddad, and H. Naira. Neural network adaptive control for a class of nonlinear uncertain dynamical systems with asymptotic stability guarantees. *IEEE Transactions on Neural Networks*, 19:80–90, 2008.

D. Travis, B. T. Thumati, and S. Jagannathan. Optimal control of unknown affine nonlinear discrete-time systems using offline-trained neural networks with proof of convergence. *Neural Networks*, 22:851–860, 2009.

A. Ude, C. Atkeson, and M. Riley. Programming full-body movements for humanoid robots by observation. *Robotics and Autonomous Systems*, 47:93–108, 2004.

S. Wang, H. Luo, C. Yue, and X. Liao. Parameter identification of chaos system based on unknown parameter observer. *Physics letters. A*, 372:2603–2607, 2008.

H. Wei and S. Amari. Dynamics of learning near singularities in radial basis function networks. *Neural Networks*, 21: 981–1005, 2008.

N. Xie and H. Leung. Blind identification of autoregressive system using chaos. *IEEE Transactions on Circuits and Systems*, 52:1953–1965, 2005.

K. Yamane, J. J. Kuffner, and J. K. Hodgins. Synthesizing animations of human manipulation tasks. *ACM Trans. Graph.*, 23(3):532–539, 2004. ISSN 0730-0301. doi: http://doi.acm.org/10.1145/1015706.1015756.

K. Yokoi, E. Yoshida, and H. Sanada. Unified motion planning of passing under obstacles with humanoid robots. In *Robotics and Automation, 2009. ICRA '09. IEEE International Conference on*, pages 1185–1190, May 2009. doi: 10.1109/ROBOT.2009.5152797.

E. Yoshida, M. Poirier, J. Laumond, O. Kanoun, F. Lamiraux, R. Alami, and K. Yokoi. Whole-body motion planning for pivoting based manipulation by humanoids. In *Proceedings of IEEE International Conference on Robotics and Automation*, 2008.

R. Zollner, T. Afour, and D. R. Programming by demonstration: Dual-arm manipulation tasks for humanoid robots. In *Proceedings of the International Conference on Intelligent Robots and Systems*, 2004.

///VSP///
Advanced Robotics 23 (2009) 1951–1996
ADVANCED ROBOTICS
www.brill.nl/ar

Full paper

Effects of Embodiment and Gestures on Social Interaction in Drumming Games with a Humanoid Robot

Hatice Kose-Bagci *, **Ester Ferrari**, **Kerstin Dautenhahn**, **Dag Sverre Syrdal and Chrystopher L. Nehaniv**

Adaptive Systems Research Group, University of Hertfordshire, School of Computer Science, Hatfield, AL 10 9AB, UK

Received 9 June 2009; accepted 8 July 2009

Abstract

We present results from an empirical study investigating the effect of embodiment and minimal gestures in an interactive drumming game consisting of an autonomous child-sized humanoid robot (KASPAR) playing with child participants. In this study, each participant played three games with a humanoid robot that played a drum whilst simultaneously making (or not making) head gestures. The three games included the participant interacting with the real robot (physical embodiment condition), interacting with a hidden robot when only the sound of the robot is heard (disembodiment condition; note that the term 'disembodiment' is used in this paper specifically to refer to an experimental condition where a physical robot produces the sound cues, but is not visible to the participants), or interacting with a real-time image of the robot (virtual embodiment condition). We used a mixed design where repeated measures were used to evaluate embodiment effects and independent-groups measures were used to study the gestures effects. Data from the implementation of a human–robot interaction experiment with 66 children are presented, and statistically analyzed in terms of participants' subjective experiences and drumming performance of the human–robot pair. The subjective experiences showed significant differences for the different embodiment conditions when gestures were used in terms of enjoyment of the game, and perceived intelligence and appearance of the robot. The drumming performance also differed significantly within the embodiment conditions and the presence of gestures increased these differences significantly. The presence of a physical, embodied robot enabled more interaction, better drumming and turn-taking, as well as enjoyment of the interaction, especially when the robot used gestures.

Keywords

Human–robot interaction, embodiment, gestures, humanoid, drumming game, social interaction

* To whom correspondence should be addressed. E-mail: hatice.kose@gmail.com

 DOI:10.1163/016918609X12518783330360

1. Introduction

Social robots are being used widely in interaction games with human interaction partners, e.g., in the application areas of entertainment robotics [1–4], socially assistive robotics [5] and robot-assisted therapy [6–9]. Their physical appearance as well as behavior affect the participants and motivate them to take part in the interaction. The use of social cues such as gestures produced by the robots also has a great impact on the motivation of the human participants. In human–human interaction, gestures play an important role in communication, coordination and regulation of joint activities. In the related field of virtual agents, the beneficial effects of gestures and expressions used by virtual agents were shown both in short-term and long-term interactions, in maintaining user involvement with the tasks encouraged by the agent [10, 11].

Given that the term ‘social robots’ refers to robots that are designed to evoke meaningful social interaction with their users [12], social robots may not necessarily need a physical body to accomplish their goals [13], unlike other robotic systems where the task requires a physical form (e.g., object manipulation). Socially interactive robots are used in various kinds of applications, such as toys for entertainment, rehabilitation aids or educational tools. Often their primary functionality is not strictly related to physical interaction or manipulation, which implies that a physical body may not be required (a virtual embodiment could be sufficient). For example, in the field of rehabilitation robotics, which often focuses on physically assistive robots, non-contact socially assistive robots have been developed with the primary goal to motivate and monitor the user during the rehabilitation phase [14, 15]. Thus, this raises the question of whether physical embodiment is essential for successful interactions between a human and a social robot.

In the context of this paper we follow the notion of embodiment that has been defined as “that which establishes a basis for structural coupling by creating the potential for mutual perturbation between system and environment” [16, 17]. Note that this definition of embodiment does not necessarily require a system to possess a physical shape. According to Chrisley and Ziemke [18], embodiment can be classified in four different levels, from physical realization, where the system must simply be realised in some physical substrate, to organismal embodiment where the body must be alive (i.e., metabolize, reproduce, etc.).

Previous research has shown that physical embodiment has positive effects on the quality of interaction between social robots and humans. Lee *et al.* [19] conducted two experiments to investigate the effects of physical embodiment and tactile communication in human–agent interaction. They found that physical embodiment, as a bodily presence, played an important role in social interactions between human and social agents, although social robots were not particularly related to physical functions. Participants preferred interactions with physical social robots to interaction with virtual social robots. In the first experiment they found that physical embodiment positively impacted the agent’s social presence, as well as the evaluation of the interaction with the agent, while a second experiment showed that physical embodiment with restricted tactile interaction caused insignificant or even negative effects in human–agent interaction.

Positive effects of physical embodiment have also been found by Bartneck [20] in his study with an emotional robot (eMuu). Specifically, he claimed that physical embodiment facilitated social interaction. In an empirical study, participants acquired a higher score in a negotiation game when they interacted with a robotic character than when they interacted with a character on a computer screen.

Tapus and Mataric [21] studied the effect of embodiment in a human–robot interaction (HRI) experiment with social robots playing music. Here a physical robot or a simulated computer animation played recorded songs and patients with cognitive impairments tried to distinguish between up to four songs. Future work needs to demonstrate whether the physical embodiment of the robot motivated the patients positively and helped them to improve their cognitive impairments.

Two experiments conducted by Wainer *et al.* [15, 22] add further support to the importance of physical embodiment on performance and impression of social interactions. Their results demonstrate that a physically situated robot is more appealing than a non-embodied robot, and it is also seen as more helpful, watchful and enjoyable when compared to a remotely tele-present robot and a simulated one. These results suggest that a physical robot may be more effective in assistive physiotherapy (the particular application area that the authors aimed at) than a disembodied one.

Conversely, a study conducted by Powers *et al.* [23] showed that an interaction with a ‘collocated robot’ (physical embodied robot) compared to a remote projected robot does not always lead to better results. They found that the projected robot had almost as much social influence as the collocated one, i.e., it was equally engaging, elicited equal disclosure, but may have had somewhat less influence (the participants did not rate the projected robot as highly when they evaluated its helpfulness, the usefulness of its advice and its effectiveness as a communicator).

Interestingly, all the above results are related to studies that have been conducted with adults. Children, even if they are the main target in the area of entertainment robotics, have not been involved extensively in such research comparing physical and virtual embodiment conditions. One of the few studies conducted with children participants has been carried out by Pereira *et al.* [2], involving 18 children in a gaming scenario against a robotic agent or a virtually embodied agent. In spite of the results of this study suggesting that embodiment has a positive effect on participant’s enjoyment, there is still sparse evidence of the effect that embodiment has on children and further investigation is needed.

A social robot needs a set of social skills in order to successfully encourage a user’s social behavior, which might require the ability to use social cues and gestures to motivate users to interact with it and keep them motivated to interact with the robot beyond the first few moments of ‘novelty’. This is especially the case for assistive robotics [5]. A variety of robotic systems have been using social cues and gestures in order to encourage HRI, e.g., KISMET [24, 25], where the interaction itself was the primary goal. Different from this work, our own studies include an enjoyable task that will need to be achieved jointly by the human–robot pair. Therefore, we have chosen drumming as a test bed for our studies. Drumming is relatively straightforward to implement and test, and can be implemented technically without special actuators like fingers or special skills. Additionally, it is an easy and enjoyable task for the participants, who do not require any detailed information or skills. There are several approaches concerning drumming in HRI. For example, robotic percussionists play drums in collaboration with human interaction partners, where they use robotic arms that are specially designed to play drums [26, 27]. Similarly, humanoid drumming is used as a test bed for exploring synchronization [28].

Drum-mate is an interaction-based imitation game based on the autonomous drumming game of a human interaction partner and a humanoid robot [29, 30]. In Drum-mate studies, the humanoid robot KASPAR (Kinesics and Synchronization in Personal Assistant Robotics) plays drums autonomously with a human 'partner' (interactant), trying to imitate the rhythms produced by the human. However, the social interaction is not limited to the replication of drumming, but also involves studying the impact of non-verbal robot gestures that are meant to motivate the human. KASPAR produces head gestures from a limited repertoire and eye-blinking as it drums. KASPAR is a minimally expressive child-sized humanoid robot developed previously by our research group for its use in human–robot social interaction games (for more details of the robot, see Ref. [31]). In previous work, two studies with 24 adult participants each [29, 30] analyzed the interaction between participants and the humanoid robot in terms of imitation, turn-taking and the impact of non-verbal gestures as social cues [30]. Different computational probabilistic models were used to achieve turn-taking that is not deterministic, but emerging from the interaction between the human and the humanoid robot. The humanoid robot is no longer a passive 'follower', but can also play the 'leader' role in the game [29]. Also, different orders of the game conditions were tested and a significant effect of play time was found. The error rates in drumming and turn-taking significantly decrease as the human players play more games [29, 30].

The above-mentioned Drum-mate scenario (with adult participants) formed the basis for the current study which is a modified version of the Drum-mate game. It was tested with 66 primary school students, where different embodiment conditions, together with their relation with the head gestures, were studied. Each participant played three interactive drumming games with the humanoid robot. In each game the participant interacted either with the real robot (physical embodiment condition), with a real-time image of the robot (virtual embodiment condition) or with the hidden robot (disembodiment condition) — in this last condition only the sound of the robot is heard. Half of the children played games while the humanoid robot was simply drumming without making any head gestures; during the games with the other half of the children, KASPAR played its drum whilst simultaneously making gestures and waving its hand 'good-bye' at the end of the game.

Compared to the previous Drum-mate experiments with adults, several modifications were made in the current study to adapt the experiment to the child participants (e.g., simpler gestures were used, and the single game duration and the time between turns were decreased).

The paper is organized as follows. Section 2 presents the research questions and hypotheses. The experiment design and data collection are described in Section 3. In Section 4, the experimental results are described. Section 5 includes a brief conclusion of the experiment, lessons learned and presents ideas for future work.

2. Research Questions and Hypotheses

The goal of the study was to determine whether embodiment and gestures have an effect on how users perceive a social robot. In this experiment, we examined three levels of embodiment, each of which was used both with and without robotic gestures.

To study the effects of the embodiment, each child played a drumming game in the following three conditions:

- *K* The physical embodiment condition in which KASPAR sat on a table and played a drum in front of the child;
- *V* The virtual embodiment condition in which KASPAR's image was projected on the wall (keeping its real dimension in terms of size), while the robot played a drum behind an opaque barrier that separated it from the child;
- *D* The disembodiment condition (only sound), in which KASPAR and the child were in two different areas, and KASPAR played a drum behind an opaque barrier. Note that in this case participants were not able to see the robot, but they could actually hear when it was producing gestures (i.e., in the gesture condition the children could hear the robot's motors moving behind the screen).

Note that different areas had to be created for practical reasons, i.e., in order to allow each child to be tested in three experimental conditions and allowing quick changes between the experimental settings.

Based on the results of previous research (see Section 1), we expected that a social robot, in order to be able to engage in a playful interaction with a child, would require a certain degree of embodiment. Thus, we investigated the following hypotheses:

- H1 Children would evaluate a social robot and the interaction with it more positively when they played with an embodied robot (conditions K and V), than when they interact with a disembodied robot (condition D); and comparing the physical and virtual embodiment conditions, they will evaluate the physically embodied social robot (K) more positively. Moreover, we expected that the presence of gesture would increase the difference in how the children evaluate the embodiment conditions. Specifically, we expected the children would evaluate the case K (physical embodiment) more positively than the other cases (V and D) when gestures were used.

H2 The error rates in drumming and turn-taking would decrease when the children played with an embodied robot (conditions K and V) compared when they interact with a disembodied robot (condition D); and comparing the physical and virtual embodiment conditions, the error rates in the physically embodied robot (K) condition will be lower than the virtually embodied robot (V). The differences between the drumming and turn-taking performances when the children play with different conditions (K, V or D) are expected to be higher when the gestures were introduced.

H3 As the play time increases, the error rates would decrease. In other words, the more the child plays, the better her/his drumming and turn-taking would be. Therefore, we expect their drumming and turn-taking to improve over time and, consequently, we expect a better performance (lower error) in the third game played than the first game played. The differences between the drumming and turn-taking performances between first and third games are expected to be higher when the gestures were introduced.

To study the effects of gestures, the robot played half of the games while making a gesture ('gesture' condition) and half of the games without making any gesture ('no-gesture' condition).

In addition to the above hypotheses, we will also consider the effects of gender. Previous works show gender differences have an impact on the subjective and objective evaluation of human participants. Kose-Bagci *et al.* showed that females and males evaluate the robot and the interaction games differently, and their drumming and turn-taking performances differ significantly [30]. Here the males were more 'task oriented', whereas females tended to value interactional aspects of the scenario. However, in Ref. [32] it was found that the males like to see the robot as more 'human-like' and achieve a social facilitation, while females saw it 'machine-like'. Gender differences were also revealed in HRI experiments [33–35]. This suggests the possibility of gender as a confounding variable in this experiment, which will be examined in the data analysis below. However, gender issues do not play a major role in our research goals (and for this reason our sample is not gender balanced).

3. Experiment

3.1. Participants and Sample

Sixty-six participants in the age range of 9–10 years took part in the study. All participants were primary school students from six schools in Hertfordshire, UK. Gender was not balanced in the sample; the majority of the children were female ($n = 39$, Table 1). None of the children had interacted with the robot KASPAR prior to the experiment. Most of them were used to playing computer games (Table 2)

Table 1.
Distribution of the sample's gender

	N	%
Male	27	41
Female	39	59
	66	100

Table 2.
Distribution of the sample's familiarity with robots

	N	%
I've played with a robot before	19	29
I've never played with a robot	47	71
	66	100

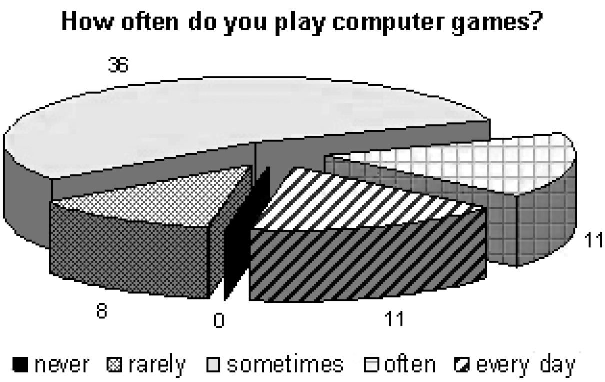

Figure 1. Distribution of how often the children play computer games.

and they were generally unfamiliar with robots (Fig. 1). Prior to the experiment the children's parents had consented for their children to take part in the study, and allowed us to obtain photos and video recordings of the experiments for scientific purposes.

3.2. Design

To test the hypotheses, a 3 (embodiment) × 2 (gestures) mixed design was used. The experiment consisted of drumming games with a humanoid robot in three different embodiment conditions (physical embodiment, virtual embodiment and disembodiment) and two gesture conditions (with and without gestures).

Each participant played an equivalent drumming game with the robot in the three different embodiment conditions, each of which took 2 min (making embodiment a repeated measures variable). Half of the children were randomly assigned to the gesture condition while the remaining half was assigned to the no-gesture condition (making gestures an independent groups variable).

In the no-gesture condition, KASPAR only played the drum without making any gestures, while in the gesture condition, the robot made simple head gestures, e.g., nodding and moving the head from side to side during its drumming session. These head gestures were played in a fixed sequence to encourage the participants to believe that they were executed on purpose and not at random [30]. KASPAR smiled when it started drumming and if it did not detect the child's drumming when expected (e.g., if the drumming beats were too light, too fast or if the child did not play), then it blinked and expressed a neutral smile. At the end of the game KASPAR waved its hand with a 'good-bye' gesture to notify the end of the game. Note that in the disembodiment condition participants did not see the gestures, but they could hear them (i.e., they could hear the robot's motors moving behind the screen) and the production of the gestures also slightly influenced the robot's behavior timing. Thus, for completeness purposes we also included the condition where the disembodied robot used gestures.

A repeated measures design has the advantage that individual differences between participants are removed as a potential confounding variable, but a drawback related to the order effects. As we assumed that the order in which the three embodiment conditions were presented could influence the participant's opinion and behavior, their order was counterbalanced and all six possible presentation orders were used. This was essential to account for possible fatigue, habituation or learning effects.

3.3. Experimental Setup

The experiment was conducted during the event called 'Take Part In The Future And FearNot!' [36], hosted by the University of Hertfordshire, School of Computer Science and School of Education in May 2008, where 8- to 12-year-old children had the opportunity to interact with robots and trial anti-bullying software. In addition, in a different large room, the children were also able to interact with a number of humanoid and non-humanoid robots that are used in our research group. While screens were used to separate the experimental area from the other robotics demonstrations, the experimental setting was challenging as it was not an easily controlled laboratory environment.

Moreover, in conducting the experiment the enjoyment of the activities for the participants was taken into account and children were encouraged to experience an interaction with a social robot in an enjoyable manner. While this setting might have made it easier for children to express a more positive opinion than they would in a different setting (all the variables' means were quite high), such a setup provided an enjoyable and relaxing environment that creates situations more similar to those where children's play naturally occurs. It has been argued by Sabanovic *et al.* [37] that 'Interactions with robots in the laboratory, under the watchful eye and expert guidance of the robot's designers, do not provide insights into the aspects of human robot interaction that emerge in the less structured real-world social settings in which they are meant to function. It is, therefore, necessary to evaluate human–robot interactions as socio-culturally constituted activities outside the laboratory, or "in the wild" ' (p. 576). Studying social and enjoyable games that children play with a robot requires a suitable setting, not one where the children may be under the impression that they are being evaluated or monitored (similar to an examination). Thus, we had to find a trade-off between such an enjoyable setting and the need to control the experiments. The solution we decided on is to have an experimental setup on University premises, but to situate it in the context of an enjoyable activity for the children. Also, our sample was an 'opportunity sample' — we could not freely select the children we worked with, which meant that we could not control all the variables (e.g., gender). Thus, our research approach is similar to other HRI experiments in public places, museums, shopping malls, school environments, etc. (e.g., Refs [37–39]).

We designed a separate experimental area in the room where the robotics demonstrations took place. Two almost identical cubicles isolated from the rest of the room by tall screens were used to carry out the study (Fig. 2). In the remainder of the room other robotic activities took place at the same time with other children. In the first cubicle (area 1), the robot KASPAR was seated on a table with a toy drum on its lap. A chair was placed in front of the robot where the participant was seated (Fig. 3). In the second cubicle (area 2), there was a table and a chair for the participant. In the virtual embodiment condition, a real-time image of KASPAR (the same size as the physical robot) was projected on the wall in front of the seated child participant. In the disembodiment condition, the projector was switched off and the child participant could just hear the drumming sound of the robot hidden behind a screen (Fig. 4). To study the effect of the different embodiment conditions, all the others features of the setup in the cubicles were kept the same.

During the experiments, a drum (with a microphone attached) and a stick to hit the drum were made available to the children to make it easier for the robot to recognize the drumming sound through audio analysis. Although some of the children used the stick to hit on the top of the drum as we suggested, others preferred to use both their hands or to hit the tambourine-style bells around the drum's sides, which increased their enjoyment and involvement in the game, but negatively effects the audio analysis. Note that since the interaction was meant to be enjoyable and playful, we did not insist on the children using the drum stick.

Participants were instructed that they could play drumming games with the robot KASPAR. Simple general instructions about the game were given (e.g., hit the drum strongly so that KASPAR can 'hear' you better).

The children entered the experimental area in groups of three. On arrival, one of the experimenters gave them a 30-s demo of the first condition they were going to play. That was done simply to show them what the game was about and to familiarize them with the setup of the cubicle (e.g., if the first condition was 'vir-

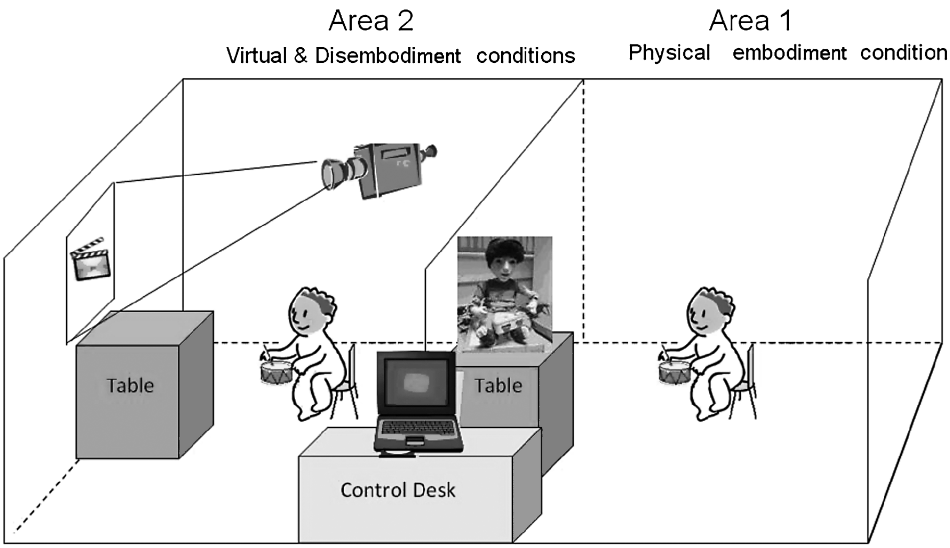

Figure 2. Setup of the experiment. The experiment took place in two identically designed cubicles. In the right cubicle (area 1), the game with the physical embodied condition was played, with KASPAR sitting on a table in front of the child. The left cubicle (area 2) was used to play the disembodiment and the virtual embodiment conditions. In the virtual embodiment condition, a real-time image (the same size as the physical robot) of the drumming robot was projected on the wall just above the desk at a height comparable to the physical embodied condition. In the disembodiment condition, the projector was switched off and only the sound of the robot hidden on the other side of the screen was presented. All the other features in the cubicles were identical. In all three conditions the child was playing alone in the cubicle. The control desk with the laptop controlling the robot was equidistant to both cubicles and the experimenter could be seen from both cubicles. Several cameras were located in the room to record the experiment.

tual embodiment', then the demo was carried out in area 2 and the children would see the experimenter playing with the projected image of KASPAR — not with the physical robot itself).

After the demo, only one child was asked to remain in the cubicle — the other two children were taken by the other experimenter outside the cubicle to wait for their turn. Each child individually interacted with the physically embodied (K), virtually embodied (V) and disembodied robot (D), in one of the two gesture conditions.

During the games, the experimenter who was located in the experimental area kept silent and did not interfere with the child's or the robot's performance. The children were given very basic information about the difference between the games (e.g., 'Now you will play with the projected image of KASPAR'; 'Now you will hear KASPAR's drumming sound, but you cannot see KASPAR itself').

The sessions were videotaped and the video recordings were used as a source of behavioral data (described in detail in the following subsections).

In the current study, KASPAR acted totally autonomously. Thus, the experimenter was always in view of the participant, especially when the physical robot

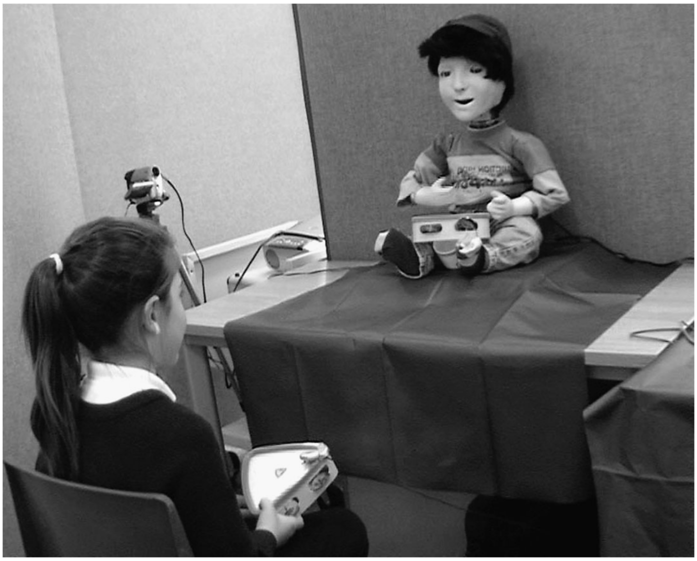

Figure 3. Screen shot from the experiment showing a person playing a drumming game with the physically embodied robot.

Figure 4. Screen shot from the experiment showing a person playing a drumming game with the hidden robot (disembodiment condition).

was not visible. Participants may otherwise have believed that the performer of the drumming was not KASPAR, but the experimenter (i.e., a hidden puppeteer in the Wizard-of-Oz technique — a widely used technique in human-computer interaction and HRI research where a human, who is unknown to the participants, is controlling the behavior of the system; e.g., Ref. [40]).

3.4. KASPAR

The study was carried out with the humanoid robot called KASPAR. KASPAR is a child-sized humanoid robot that was designed and built by the members of the Adaptive Systems Research Group at the University of Hertfordshire to study HRIs with a minimal set of expressive robot features. This humanoid robot has been used in a variety of projects, e.g., in research to mediate interaction for people with and without special needs [41, 42].

KASPAR has 8 d.o.f. in the head and neck, and 6 d.o.f. in the arms. Its width is 30 cm, depth is 35 cm and height is 45 cm, and its shape is modeled after a 2- to 3-year-old child. The face is a silicon-rubber mask, which is supported by an aluminum frame. It has eyelids capable of blinking, and a mouth capable of opening and smiling (a detailed description can be found in Ref. [31]).

3.5. Interaction Game Implementation

In this work, as in the previous study [30], the participant played a rhythm that KASPAR tried to replicate in a simple form of imitation (mirroring). KASPAR had two modes — listening and playing. It recorded and analyzed the human's rhythm in the listening mode and it played the rhythm back by hitting the drum positioned on its lap in the playing mode. Then the participant played again. This (deterministic) turn-taking in this game continued for a fixed time duration (2 min for the current work). Due to its limited motor skills, KASPAR did not imitate the strength of the beats, but only the number of beats and durations between beats. For beats beyond its motor skills, it used instead minimum values allowed by its capabilities: KASPAR needed at least 0.3 s between beats to get its joints 'ready', so that, even if the human played faster, KASPAR's imitations would still require minimum durations of at least 0.3 s between beats. It also needed to wait for a few seconds before playing any rhythm in order to get its joints into the correct reference positions.

3.6. Software Features

The implementation of robot perception and motor control used the YARP environment [43]. YARP is an open-source framework used in the project RobotCub that supports distributed computation and emphasizes robot control as well as efficiency. It enables the development of software for robots without considering a specific hardware or software environment. Portaudio [44] software was used to grab audio from the audio device within the YARP framework.

The acoustic sound waves recorded by the sound grabber module were converted to digital music samples, which allows mathematical computations and sample-based techniques to be used on them. To detect the patterns of a sound wave, a filter-based method is used, based on the work of Ref. [45] that was originally used to detect visual patterns.

3.7. Measures

During the experiment several sources were used to collect data. These sources included asking the children to complete a questionnaire related to the trials, recording the sessions by video cameras for later video analysis and collecting the behavioral (drumming) data of each robot–child pair.

3.7.1. Questionnaire

A paper-based questionnaire was used to collect data relevant in investigating the differences in the embodiment conditions (H1) and the gesture effect (H2). The questionnaire is available from E. F. request. It was comprised of three sections. The first section gathered general information about the child, their experience with robots and video games, and gave instructions on how to complete the questionnaire. The second section consisted of 15 closed-ended questions (repeated for the three embodiment conditions) and the third section consisted of two open-ended questions.

Considering the sample population and time limitations, the questionnaire was kept as simple and short as possible. It was pilot tested with one participant of a similar age group (a 9-year-old child). The pilot test confirmed that both the length of the questionnaire and the questionnaire's administration time were acceptable (below 6 min). In order to develop the final version of the questionnaire, a few minor language changes were made.

After each trial condition, the child was asked to answer the 15 closed-ended questions in order to express their opinions about the game they just played (the same questions were presented after each of the three embodiment trial conditions).

The closed-ended questions were used to evaluate the robot in terms of: enjoyment, social attraction, involvement, performance, general appearance and intelligence. As the researchers were interested in the children's feelings and opinions about the interaction with the robot, a five-point Likert scale (respondent shows the amount of agreement/disagreement with a given statement) and a semantic differential scale (a scale inscribed between two bipolar words; children select the point that most represents the direction and intensity of their feelings) were used.

Note that before conducting the data analysis, three items were removed from the questionnaire because some of the children, while they were completing the questionnaire, showed difficulties in understanding their meaning (one was a negatively phrased item).

To measure the level of enjoyment during the interaction with the robot, two questions on a five-point Likert scale with a central anchor were used ('Did you enjoy playing with this robot?' and 'Did you find it interesting?') (Cronbach's $\alpha = 0.78$). Social attraction toward the robot was measured by a modified version of McCroskey and McCain's Interpersonal Attraction Scale [46]; children were asked to indicate their level of agreement to the following statements: 'I would like to be friend with this robot' and 'I would like to spend more time with this robot' (Cronbach's $\alpha = 0.90$). Involvement in the game was measured by the level of agreement to the following statements: ‘I paid attention to the robot’ and ‘I felt that the robot involved me in the game’ (Cronbach’s $\alpha = 0.84$). The robot’s perceived intelligence was measured by participants’ levels of agreement on the statement: ‘I think this robot is intelligent’ (Cronbach’s $\alpha = 0.60$). The level of agreement was always measured using the same five-point response scales with a central anchor. Two questions concerning the robot’s performance were asked using a five-point semantic differential scale: lazy/energetic, and bad drummer/good drummer (Cronbach’s $\alpha = 0.79$). Three questions concerning the general appearance of the robot were asked using a five-point semantic differential scale: unpleasant/pleasant, not friendly/friendly and machine like/human like (Cronbach’s $\alpha = 0.78$).

Once the children had completed the items related to the third game, they were asked to complete the last section of the questionnaire. In this part they judged the overall experience by deciding which of the three games they liked the best and which they liked the least, as well as writing down the reasons behind that decision.

During the study, the questionnaire administration was performed in a dedicated area separate from playing areas 1 and 2, where one of the experimenters was on hand to help the child complete the questionnaire if needed.

3.7.2. Behavioral Data

The experiments were recorded by two different cameras positioned at different parts of the experimental area (one facing the child and the other facing the robot), during each single game. The video recordings were later analyzed manually to detect the performance of the children’s behavioral data (e.g., the number of drum beats played by the children and number of turns taken by the children at each game). This data was then compared with the behavioral data recorded by the robot itself (see below). Also, video recordings are helpful as they give valuable clues about the likes/dislikes of the children. They are also used to support the evaluation of the questionnaires.

Behavioral data that belonged to the robot and human participants were collected during the trials by the robot using its internal (joints) and external sensors (microphones). KASPAR records its performance (e.g., the number of drum beats played by KASPAR and number of turns taken by the robot during each game), as well as those of the children (e.g., the duration of time between each drum beat of the children) to imitate their performance within its physical limitations. In the next section these recordings will be described and analyzed in detail.

The behavioral data includes several parameters related to the children’s and to KASPAR’s drumming, i.e., the number of turns in a specific game, total, average and maximum number of drum beats performed by human participants and KASPAR per turn, and the drumming and turn-taking errors. The drumming error is the difference between KASPAR’s actual drumming, i.e., the number of beats KASPAR plays in a particular turn, and the number of beats the child plays. Likewise, the turn-taking error is based on the difference between KASPAR’s and the children’s turn-taking. Thus, the drumming and turn-taking errors reflect the discrepancy between human and robot drumming performance in this imitation game.

Although the robot's performance is the same under all conditions, the child's response to the robot's play differs. This affects the robot's detection and imitation of the child's drumming and, thus, influences the robot's performance in its response. There are many different reasons for the robot's erroneous detection of the child's performance (number of drumming beats and turns) caused by the children. For example, they will beat the drum while the robot is drumming and not listening (improper synchronization), so their beating will not be considered by the robot. Likewise, they may beat very fast or very light, which will not be detected by the robot, or they will use the bells of the drum, resulting in the robot detecting more than one beat.

In Section 4, several 'error' and performance measures based on the behavioral data are used to analyze the differences between different conditions. The error and performance measures are either presented per game or per turn. A game comprises the whole interaction occurring within one embodiment condition in a limited time period (2 min as specified in the current work) including several turns. The term non-zero turns defines the turns where at least one drum beat is played. For clarity all the zero turns (i.e., turns where no drum beat is played/or detected) were removed from the data for the following analysis. The term *Diffsum* (1) stands for the difference of the total sum of beats between participant and robot per game. The maximum number of beats per game shows the maximum number of drum beats played in a single turn per game. As shown in (2), *Errorsum* is *Diffsum* per number of non-zero turns (maximum of human and KASPAR). The term non-perfect turns is used for the number of non-zero turns where the number of drum beats in both the human's and KASPAR's turns do not match. If the number of non-zero turns of both do not match then the difference is also counted as a non-perfect turn. However, due to errors in observations and differences in children's play rhythms, this measure can be erroneous, giving a higher error rate than the real case, so we also take *Errorsum* and other performance measures into consideration. The term *Errorturn* defines the number of non-perfect turns per number of non-zero turns (maximum of human and KASPAR) (3).

$$Diffsum = \sum \mathrm{Beats}_{\mathrm{Human}} - \sum \mathrm{Beats}_{\mathrm{KASPAR}} \tag{1}$$

$$Errorsum = \frac{Diffsum}{\max(\text{non-zero\_turns}_{\mathrm{Human}}, \text{non-zero\_turns}_{\mathrm{KASPAR}})} \tag{2}$$

$$Errorturn\colon \frac{\text{non-perfect\_turns}}{\max(\text{non-zero\_turns}_{\mathrm{Human}}, \text{non-zero\_turns}_{\mathrm{KASPAR}})}. \tag{3}$$

To evaluate the success of a performance, the error rates, especially *Errorsum* and *Errorturn*, are taken into consideration. The lower the *Errorsum* and *Diffsum*, the better the drumming. Similarly, lower *Errorturn*, difference of non-zero turns or number of non-perfect turns values indicate better turn-taking. Ideally, *Errorsum* and *Errorturn* should be smaller than 1, and as close to 0 as much as possible. Other criteria, e.g., number of non-zero turns, average or maximum number of beats played per turn and number of beats played per game, can differ according to different conditions in the game or different features of the participants. For example, a higher average number of beats per turn might indicate more involvement of the human participant in that particular game compared to the other conditions, even though this might increase the errors in the game, due to the technical limitations of KASPAR's audio capture.

4. Results and Discussion

As mentioned above, the present study utilized a 3 (embodiment) $\times$ 2 (gestures) mixed design to evaluate embodiment and gesture effects. Data of the 66 children were analyzed using SPSS software (version 16 for Windows) and results are reported below. Detailed information about the descriptive data related to the questionnaires and behavioural data is listed in Appendix A.

In the following subsections the letter N stands for 'no-gesture' condition and the letter G stands for 'gesture' condition, i.e., they indicate whether the gestures of KASPAR were used in that particular game or not. Likewise, as explained above for the embodiment conditions, K stands for the game where the human participant played with the physical robot KASPAR, V is the virtual embodiment condition where the human played with the projected image of KASPAR and D is the disembodiment condition where the participant cannot see the robot, but can only hear the sound of the hidden KASPAR.

There are two main sources of error — the differences in KASPAR's and the child's drumming, and the differences in their turn-taking. The video recordings and the data recorded by the robot itself were analyzed to obtain these error and performance measures, which are very useful to study the behaviors of the robot and the children, and detect some significant differences between various conditions that are hard to detect from the questionnaire data only.

4.1. Does Embodiment Matter?

Data collected from the questionnaires were analyzed to investigate differences in children's opinions about their interaction with the robot in the different embodiment conditions.

In our research we predicted that the children would evaluate the interaction with an embodied robot more positively than the interaction with a disembodied one and that the interaction with the virtual robot would be less positive than the interaction with the physical embodied robot (as described in H1).

Hypothesis H1 is partially supported by the answers that the children gave to the overall experience. Two questions at the end of the questionnaire focused on collecting information about the game that they liked the most (Fig. 5) and the game that they liked the least (Fig. 6). As shown in Table 3, almost all the participants ($n = 55$; 83.33%) preferred to play with the embodied robot (conditions V or K), rather than with the disembodied one. Likewise, more than half of the children preferred the game in the physical embodiment condition ($n = 38$; 57.57%), compared

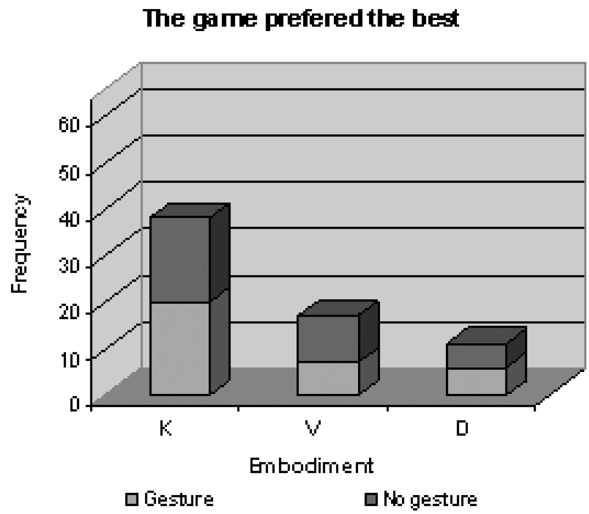

Figure 5. Most liked game.

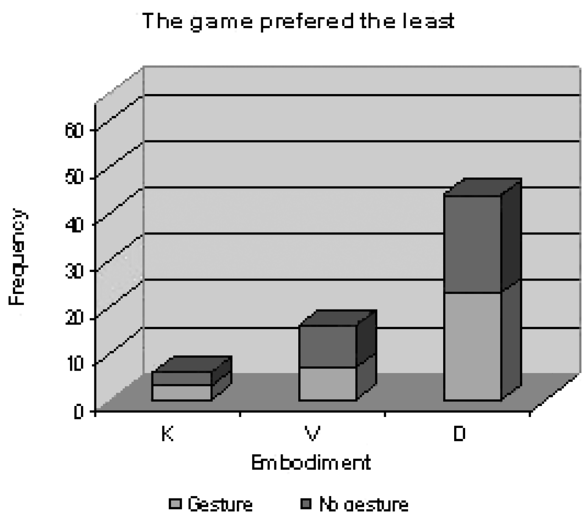

Figure 6. Least liked game.

to the game in the virtual embodiment condition (25.76%). In addition, they liked least the game played in the disembodiment condition ($n = 44$, 66.6%). However, no significant mean difference in the embodiment condition related to evaluation of enjoyment, social attraction, involvement, performance, general appearance and intelligence was found.

It is nevertheless interesting to note the trend that the data shows (Fig. 7). The interaction with the physically embodied robot was generally more appreciated than the interactions in the other two conditions. Indeed, in almost all the variables (e.g., robot's appearance, social attraction, involvement and intelligence) the

Table 3.
Game preferences

Embodiment	Game liked the best				Game liked the least			
	Frequency	Frequency in G condition	Frequency in N condition	%	Frequency	Frequency in G condition	Frequency in N condition	%
K	38	20	18	57.57	6	3	3	9.09
V	17	7	10	25.76	16	7	9	24.24
D	11	6	5	16.67	44	23	21	66.67
Total	66			100	66			100

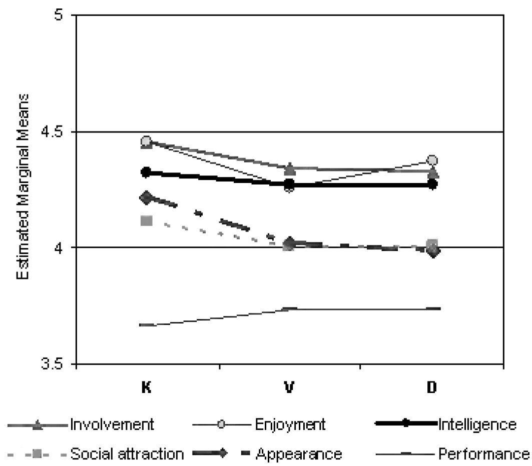

Figure 7. Effect of embodiment in children's scores of robot appearance, social attraction, involvement, intelligence, performance and enjoyment of the interaction.

children gave the physically embodied robot the highest score, the disembodied robot the lowest score and the virtual embodied robot a score in between the two (generally close to the one assigned to the disembodiment condition). On the contrary, a different result appeared for the level of enjoyment. It is interesting to notice that the virtual embodiment condition received the lowest score — lower than the disembodied condition. Similarly, another result that does not follow the previous highlighted trend is the one related to the robot's performance. The disembodiment condition gained the highest score in robot performance, while the physically embodied robot received the lowest score. A possible explanation for this unforeseen result is that children's attention, while interacting with the disembodied robot, is not diverted from the primary task, so children were focused only on the drumming game and, thus, evaluated the drumming performance differently.

However, consistent with hypothesis H1, results show that when gestures were used, the participants tended to evaluate the physically embodied robot more pos-

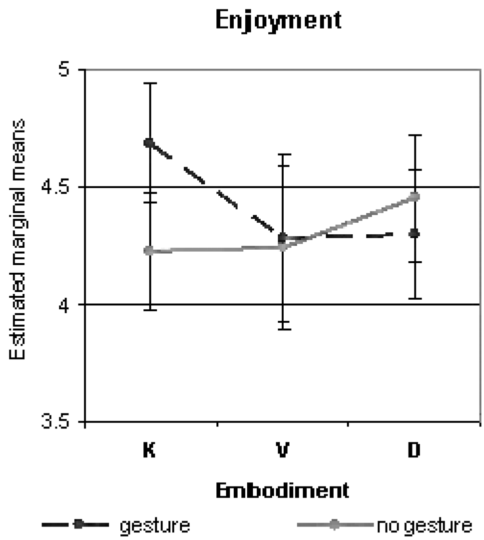

Figure 8. Combined effect of embodiment and gesture in children's scores of the enjoyment of the interaction.

itively than in the other two conditions. In terms of enjoyment, there was an interaction effect between gesture and embodiment ($F(2,126) = 4.29$, $P < 0.016$, $\eta^2 = 0.064$). This effect is described in Fig. 8, which suggests that for the gesture condition, participants tended to evaluate the physical condition more favorably than other conditions — an effect which is not evident for the no-gesture condition.

Likewise, there was an interaction effect between gestures and embodiment in terms of perceived intelligence ($F(2,126) = 3.24$, $P < 0.042$, $\eta^2 = 0.049$). This effect is described in Fig. 9, which suggests that the physical robot is rated as more intelligent in the gesture condition while the opposite is true for the virtual embodiment and the disembodied condition. It may be that the time the robot spent in making gesture movements affected negatively the intelligence attributed to it when it was not physically present (V and D condition), while it could have had a positive effect for the robot in the K condition.

It is also interesting to notice that the disembodied robot (D) in the no-gesture condition and the physically embodied robot (K) in the gesture condition received a similar score. This result might be related to the number of drum beats played per turn (see Section 4.2).

Moreover, a significant interaction effect between gestures and embodiment was found in terms of robot appearance ($F(2,128) = 4.92$, $P < 0.009$, $\eta^2 = 0.071$). Note, as mentioned, children neither saw the robot's appearance nor its gestures during the disembodiment condition, so we refrain from discussing in more detail any results concerning the disembodiment condition with respect to appearance evaluation or effects of gestures. Figure 10 suggests that this effect caused the children to evaluate the robot's appearance in the virtual condition more positively when ges-

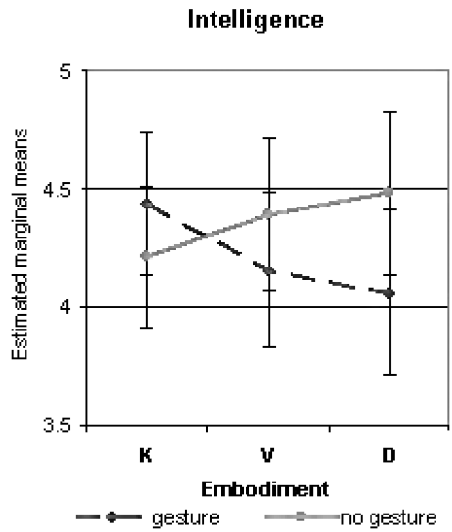

Figure 9. Combined effect of embodiment and gesture in robot's intelligence.

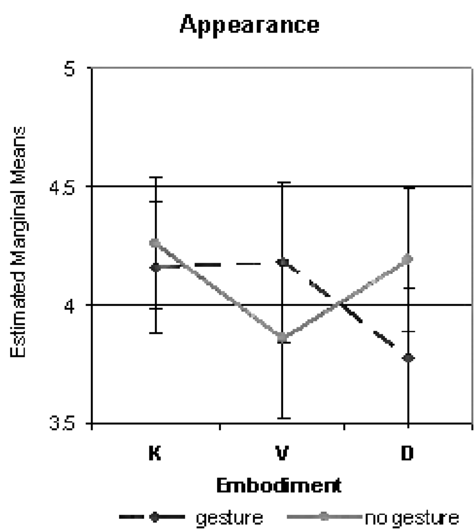

Figure 10. Combined effect of embodiment and gesture in robot appearance.

tures were used than when they were not used. Different from the previous trend, children assigned a higher score to robot appearance when the physically embodied robot was drumming in the no-gesture condition than in the gesture condition. It might be the case that the gestures made by the robot were appropriate for a projected image of the robot, but not smooth enough for a robot sat in front of the participants.

No significant interaction effects between the gesture conditions and embodiment conditions concerning social attraction, performance and involvement were found.

4.2. Effect of Embodiment and Gestures on Behavioral Data (H2)

The drumming and turn-taking performances differ for different embodiment conditions based on the analysis of behavioral data including error rates on the drumming and turn-taking, and several other parameters, i.e., total amount of drumming per game and maximum number of drum beats per turn, as stated in Section 3.7.2.

The game types are compared in detail in Table 4 (differences between human and KASPAR's perspectives), Table 5 (human's perspective) and Table 6 (KASPAR's perspective). As shown in Fig. 11, there is a highly significant difference between the physically embodied condition K and the other conditions

Table 4.
Observed differences between child–robot drumming behaviors according to the embodiment condition

Game type	Difference of sum of beats	Difference of non-zero turns	No. of non-perfect turns	Error in sum of beats	Error in number of turns
K	12.53 ± 11.0	1.86 ± 2.00	14.65 ± 4.1	0.62 ± 0.51	0.74 ± 0.13
V	18.38 ± 13.2	2.65 ± 3.40	14.98 ± 4.1	0.94 ± 0.66	0.78 ± 0.14
D	18.20 ± 16.5	2.47 ± 2.39	15.45 ± 4.0	0.90 ± 0.80	0.78 ± 0.14

Table 5.
Observed human drumming behavior according to the embodiment condition

Game type	Sum of beats	No. of non-zero turns	Maximum no. of beats	Average no. of beats/turn
K	60.06 ± 4.53	19.45 ± 4.02	12	3.18 ± 0.89
V	67.73 ± 17.72	18.70 ± 4.87	25	3.89 ± 1.73
D	67.61 ± 18.70	19.11 ± 4.18	13	3.64 ± 1.08

Table 6.
Observed KASPAR's drumming behavior according to the embodiment condition

Game type	Sum of beats	No. of non-zero turns	Maximum no. of beats	Average no. of beats/turn
K	51.86 ± 15.22	18.11 ± 3.57	14	2.93 ± 0.97
V	54.68 ± 17.81	17.59 ± 4.22	23	3.31 ± 1.73
D	53.35 ± 17.51	18.21 ± 3.52	13	2.98 ± 1.10

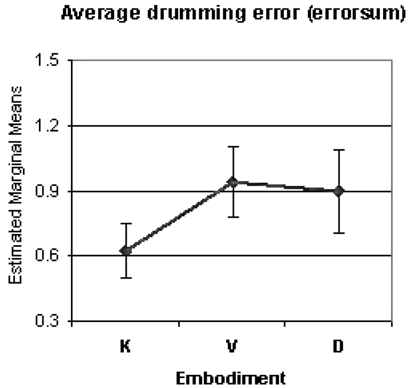

Figure 11. Average errors in sum of beats according to the embodiment condition.

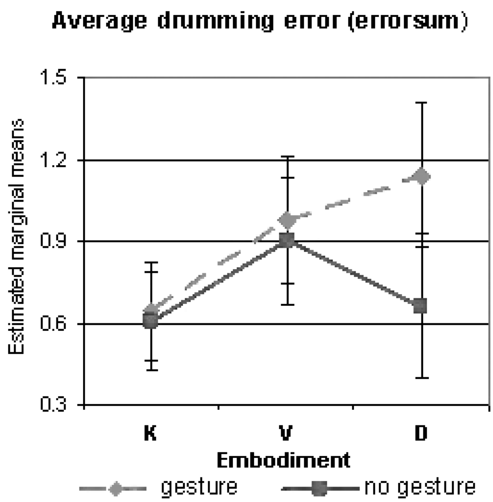

Figure 12. Effect of gestures on the average errors in the sum of beats according to the embodiment condition.

in terms of the average error in the sum of beats (*Errorsum*) ($F(2,126) = 6.653$, $P < 0.002$, $\eta^2 = 0.094$; Bonferroni adjusted *post-hoc* tests indicate that the mean difference between K and V is 0.315, $P = 0.047$; the mean difference between K and D is 0.276, $P = 0.001$). The error rate increases with the absence of physical embodiment, but not between disembodiment and virtual embodiment conditions. There is a significant interaction effect between the gesture condition and embodiment condition ($F(2,126) = 3.379$, $P < 0.037$, $\eta^2 = 0.050$). This interaction effect, presented in Fig. 12, suggests that the differences found between condition K and the others is more pronounced for the gesture condition. There is no significant main effect between embodiment conditions for turn-taking error (see Fig. 13), but a significant interaction effect ($F(2,126) = 8.214$, $P < 0.000$, $\eta^2 = 0.114$) presented in Fig. 14 suggests that gestures decrease the likelihood of such errors in

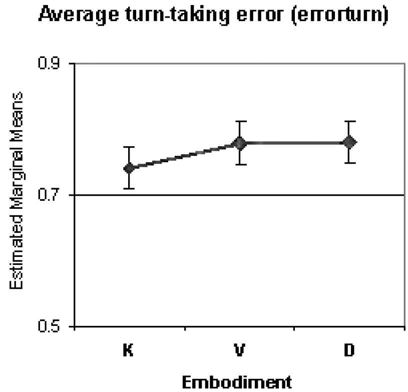

Figure 13. Average errors in number of turns according to the embodiment condition.

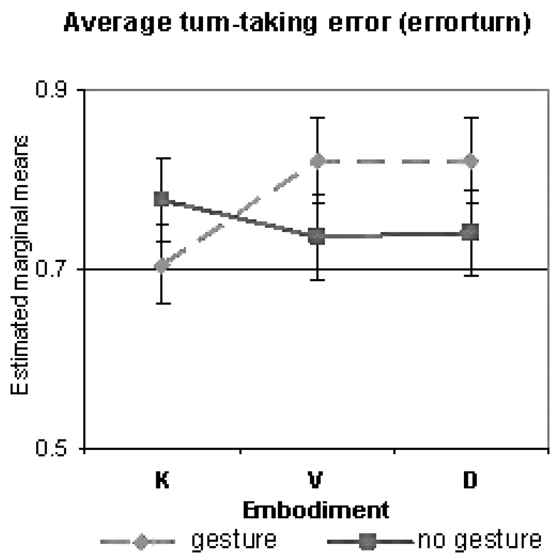

Figure 14. Effect of gestures on the average errors in number of turns according to the embodiment condition.

the physically embodied condition, while having the opposite effect on the virtual and disembodied condition. These results support hypothesis H2 and suggest that the physically embodied robot helps the child to understand the game and the robot better, and improve the drumming and turn-taking performances.

The difference in the sum of drum beats is also significantly higher in the V and D conditions than in condition K (main effect: $F(2,126) = 6.096$, $P < 0.003$, $\eta^2 = 0.087$; Bonferroni adjusted *post-hoc* tests indicate a mean difference between K and V of 8.939, $P = 0.0.001$; the mean difference between K and D is 7.591, $P = 0.0001$), while there are no significant differences between the V and D conditions (Fig. 15). An interaction effect between gesture and embodiment conditions,

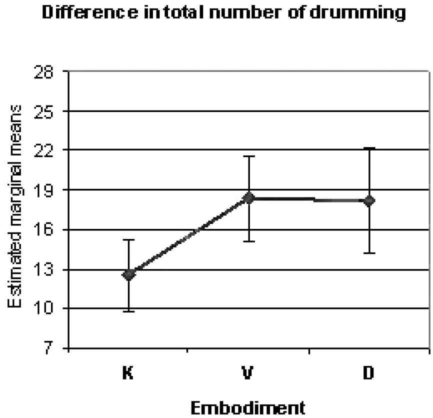

Figure 15. Difference of total number of drum beats of children and KASPAR according to the embodiment condition.

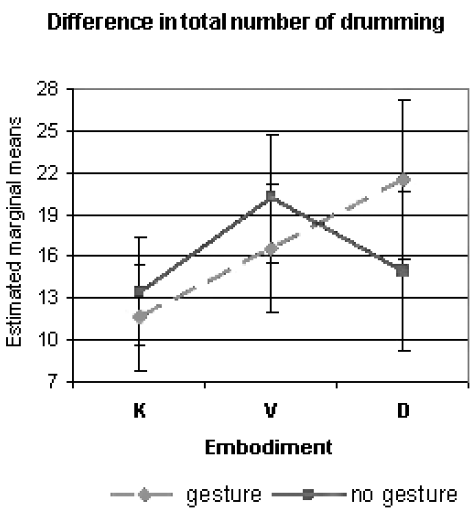

Figure 16. Effect of gestures on the difference of total number of drum beats of children and KASPAR according to the embodiment condition.

presented in Fig. 16, suggests that this effect is more strongly pronounced for the gesture condition ($F(2,126) = 4.031$, $P < 0.020$, $\eta^2 = 0.059$) (Fig. 16).

The maximum number of drum beats ($F(2,126) = 14.095$, $P < 0.000$, $\eta^2 = 0.180$; Bonferroni adjusted *post-hoc* tests indicate that the mean difference between K and V is 2.015, $p = 0.0001$; the mean difference between K and D is 1.076, $P = 0.001$) (Fig. 17), and the average number of drum beats ($F(2,126) = 13.442$, $P < 0.000$, $\eta^2 = 0.174$; Bonferroni adjusted *post-hoc* tests indicate that mean difference between K and V is 0.802, $P = 0.0001$; the mean difference between K and D is 0.495, $P = 0.0001$) (Fig. 18) of children and KASPAR per game is significantly lower in condition K than it is in conditions V and D. There is also an interaction

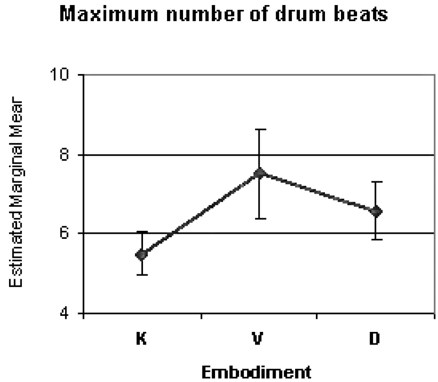

Figure 17. Maximum number of drum beats of children according to the embodiment condition.

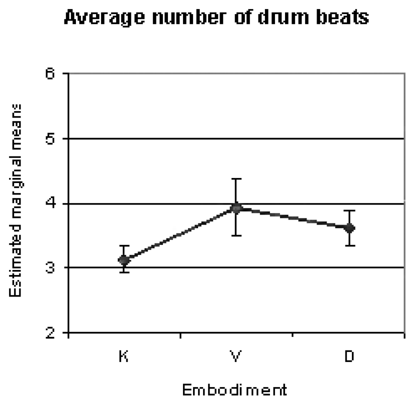

Figure 18. Average number of drum beats of KASPAR according to the embodiment.

effect that suggests that these differences are more pronounced in the presence of gestures, described in Fig. 19 ($F(2,126) = 5.145$, $P < 0.007$, $\eta^2 = 0.074$).

4.3. *Effect of the Play Time on the Games and Gestures (H3)*

The analysis of the behavioral data showed significant effects for the play time in terms of the drumming and turn-taking performances. In general, it was observed that the participants initially either tried very long and fast patterns or they did not beat the drum loud enough to be detected reliably when they started to play (KASPAR uses a high-level noise filter to filter out high inner noise coming from its joints, so it can only sense loud beats).

The effect of the play time is described in detail in Table 7 (differences between human and KASPAR's perspectives), Table 8 (human's perspective) and Table 9 (KASPAR's perspective). Interestingly, without any external encouragement, as the children played more, it appeared that they got used to the game and were progressively able to synchronize themselves to the robot better. Thus, the error rate (*Errorsum*) decreased significantly over time (Fig. 20) ($F(2,126) = 7.563$, $P < 0.001$,

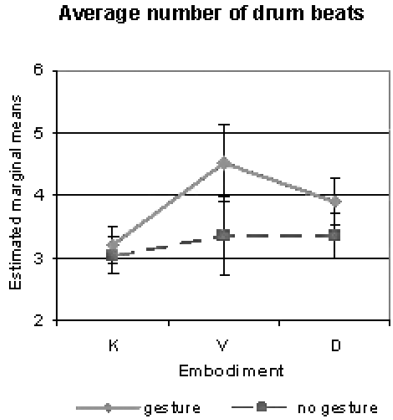

Figure 19. Effect of gestures on the average number of drum beats of humans according to the embodiment condition.

Table 7.
Observed differences between KASPAR-human drumming behaviors according to game order

Order	Difference of sum of beats	Difference of non-zero turns	No. of non-perfect turns	Error in sum of beats	Error in number of turns
1	20.15 ± 16.6	2.53 ± 2.3	15.58 ± 3.8	1.00 ± 0.79	0.79 ± 0.11
2	16.71 ± 13.2	2.64 ± 3.3	15.38 ± 3.7	0.82 ± 0.63	0.77 ± 0.14
3	12.24 ± 10.4	1.82 ± 2.2	14.14 ± 4.6	0.64 ± 0.55	0.74 ± 0.16

Table 8.
Observed human drumming behavior according to game order

Order	Sum of beats	No. of non-zero turns	Maximum no. of beats	Average no. of beats/turn
1	64.15 ± 19.2	19.15 ± 4.57	12	3.44 ± 0.99
2	64.92 ± 16.2	19.56 ± 4.20	13	3.44 ± 1.08
3	66.32 ± 16.8	18.55 ± 4.32	25	3.83 ± 1.71

Table 9.
Observed KASPAR's drumming behavior according to game order

Order	Sum of beats	No. of non-zero turns	Maximum no. of beats	Average no. of beats/turn
1	47.52 ± 15.9	17.86 ± 3.53	14	2.70 ± 0.98
2	52.33 ± 16.5	18.02 ± 3.77	13	2.96 ± 1.07
3	60.05 ± 15.9	18.03 ± 4.06	23	3.56 ± 1.66

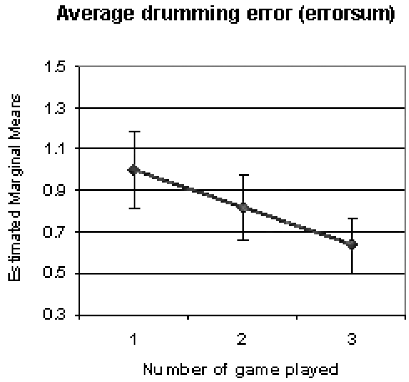

Figure 20. Average errors in sum of beats according to order.

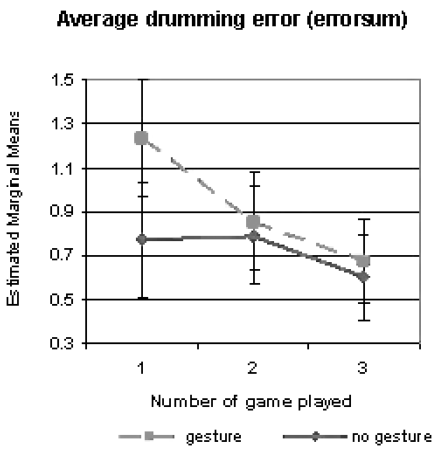

Figure 21. Effect of gestures on the average errors in sum of beats according to order.

$\eta^2 = 0.106$; Bonferroni adjusted *post-hoc* tests indicate that the mean difference between the first game and the third game is 0.365, $P = 0.003$), implying a significant interaction effect between order and gestures, suggesting that this effect is more pronounced in the gesture condition. ($F(2,126) = 2.952$, $P < 0.056$, $\eta^2 = 0.044$), as shown in Fig. 21. Similarly, the difference between KASPAR's and human's total drumming decreases as the children play more games ($F(2,126) = 7.067$; $P < 0.001$, $\eta^2 = 0.099$; Bonferroni adjusted *post-hoc* tests indicate that the mean difference between the first game and the third game is 7.242, $P = 0.005$; the mean difference between the second and the third game is 5.106, $P = 0.008$). There is also an significant decrease of the error rate in the turn-taking (*Errorturn*) between the first and third games ($F(2,126) = 5.520$, $P < 0.022$, $\eta^2 = 0.079$) (Fig. 22) that supports hypothesis H3.

As shown in the Fig. 23, the maximum number of beats per game increased significantly between the first and third game ($F(2,126) = 7.455$, $P < 0.001$, $\eta^2 =$

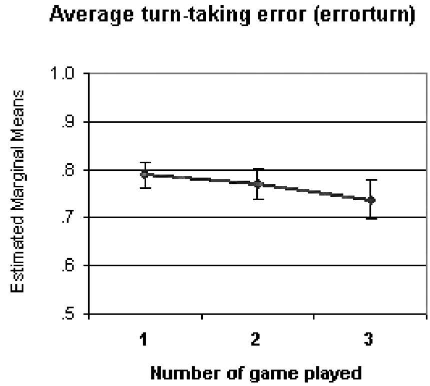

Figure 22. Average errors in number of turns according to game order.

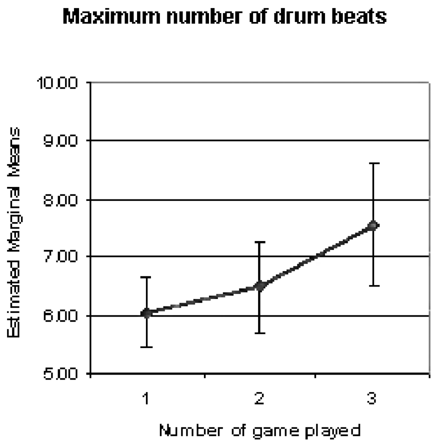

Figure 23. Maximum number of drum beats of children per turn according to order.

0.079; Bonferroni adjusted *post-hoc* tests indicate that the mean difference between the first game and the third game is 1.33, $P = 0.013$). Also, the average drumming per turn increases significantly in the third game ($F(2,126) = 4.732$, $P < 0.010$, $\eta^2 = 0.069$; Bonferroni adjusted *post-hoc* tests indicate that the mean difference between the first and the third game approaches significance, the mean difference is 0.428, $P = 0.06$) (Fig. 24), which may suggest that participants played more beats, possibly due to a stronger involvement in the game, as they played more games.

The number of non-zero turns differs significantly for the children according to order as shown in Fig. 25 ($F(2,126) = 3.800$, $P < 0.025$, $\eta^2 = 0.056$; Bonferroni adjusted *post-hoc* tests indicate that the mean difference between the second and the third game is 1.182, $P = 0.011$). Children (and consequently KASPAR) played in less turns with a higher number of beats per turn and with longer durations when the second and last games were compared.

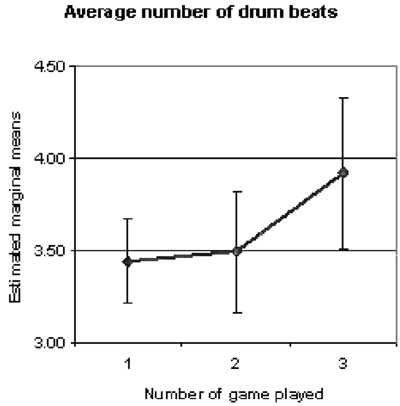

Figure 24. Average number of drum beats of children according to order.

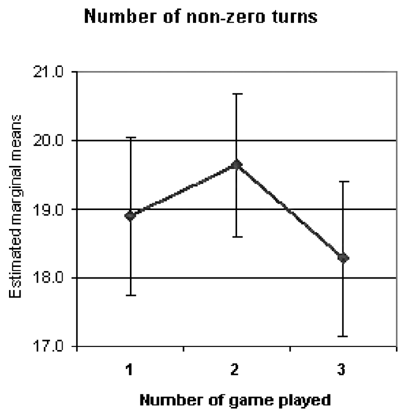

Figure 25. Number of the non-zero turns of children per game according to order.

4.4. Effect of Gender

A significant effect was found for gender in the evaluation of the robot's social attraction. This effect is described in Fig. 26 and suggests that male children found the robot more socially attractive in the physical embodiment condition than in the two other embodiment conditions, while a similar effect is not evident for females ($F(2,126) = 3,06$, $P < 0.051$, $\eta^2 = 0.046$) (Fig. 26).

In terms of behavioral data, gender showed significant differences. When different embodiment conditions were compared, there was a significant interaction effect between embodiment and gender in terms of the maximum number of beats played per turn ($F(2,126) = 8.497$, $P < 0.000$, $\eta^2 = 0.117$). The effect, shown in Fig. 27, suggest that when the children play with the physical robot, their performance is similar, but in the absence of the physical robot male children play more beats than the female children. This effect is most pronounced in the virtual embodiment condition. There could possibly be a link between the males tending to play computer games and these results. They may view the game with the two-dimensional pro-

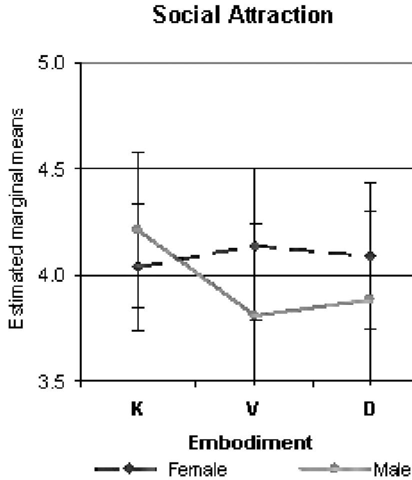

Figure 26. Gender differences in the robot's social attraction in the different embodiment conditions.

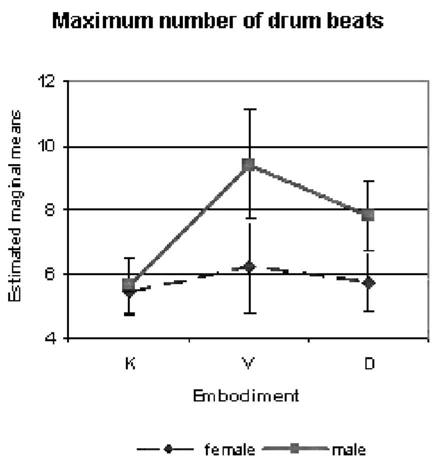

Figure 27. Effect of gender on the maximum number of drum beats of children according to the embodiment condition.

jected image of the robot (V condition) like a computer game and this might affect their performances. Further analysis is needed to investigate this issue.

In terms of the play time effect, the gender differences are significant in the number of drum beats played. There is a significant main effect for gender ($F(1,63) = 7.042$, $P = 0.01$, $\eta^2 = 0.099$; due to there only being two levels for gender, a Bonferroni test was not conducted), for game order ($F(2,126) = 7.455$, $P = 0.001$, $\eta^2 = 0.104$; Bonferroni adjusted *post-hoc* test found the following significant mean differences: first and third game, mean difference $= 1.493$, $P = 0.004$; second and third, mean difference 1.060, $P = 0.028$), as well as a significant interaction between gender and game order ($F(2,126) = 4.639$, $P < 0.011$, $\eta^2 = 0.068$). This interaction effect is described in Fig. 28 and suggests that the male participants

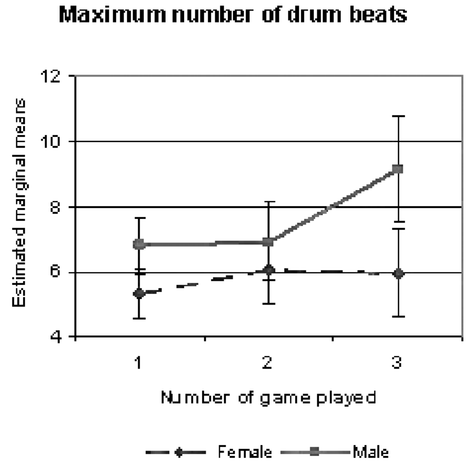

Figure 28. Effect of gender on the maximum number of drum beats of children according to the sequential order.

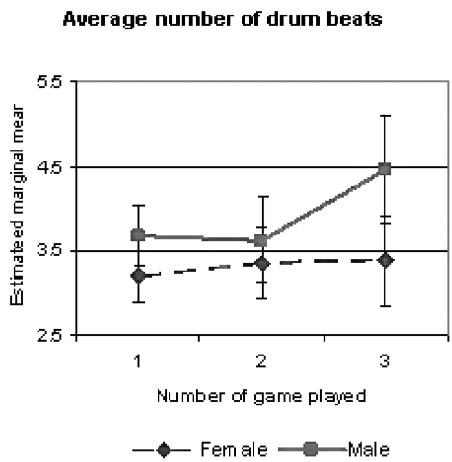

Figure 29. Effect of gender on the average number of drum beats of humans according to the sequential order.

have a more pronounced difference between the third game and the other games than the females.

For average number of beats, a significant main effect was found for gender ($F(1,63) = 5.212$, $P = 0.026$, $\eta^2 = 0.075$; Bonferroni test not conducted due to gender only having two levels) and game order ($F(2,126) = 4.732$, $P = 0.012$, $\eta^2 = 0.069$; Bonferroni adjusted *post-hoc* tests found a significant mean difference between the first and the the third game, mean difference $= 0.481$, $P = 0.035$); there was also an interaction effect between gender and game order approaching significance ($F(2,126) = 2.922$, $P < 0.057$, $\eta^2 = 0.044$). This interaction effect is described in Fig. 29, and suggests that the difference between the first and the third game is more pronounced for the male participants.

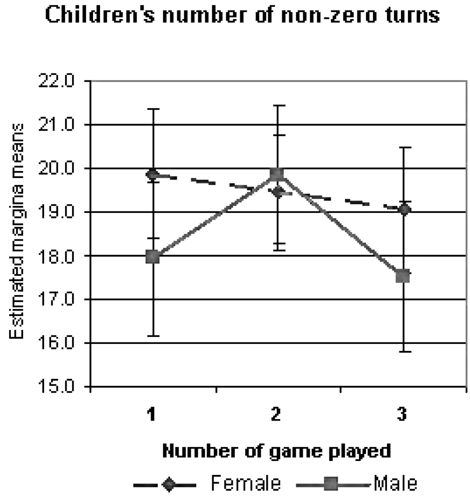

Figure 30. Effect of gender on the number of non-zero turns of humans according to game order.

For the number of turns played, a significant main effect was found for game order ($F(2,126) = 3.800$, $P = 0.025$, $\eta^2 = 0.056$; Bonferroni adjust *post-hoc* test found a significant mean difference between the the second and the third game, mean difference $= 1.359$, $P = 0.002$). However, an interaction effect was found for gender and game order ($F(2,126) = 3.263$, $P < 0.041$, $\eta^2 = 0.049$); this effect is described in Fig. 30 and suggests that this difference is primarily due to the behavior of the male participants.

5. Conclusions

In this research, we studied the effect of embodiment and gestures on a human–humanoid drumming game. We tested different levels of embodiment of the humanoid robot which autonomously played games with child participants. Half of the children interacted with a robot that made simple head gestures while imitating the child's drumming, while the rest of them played with a robot that did not make any gestures, but simply played its drum.

The analysis of results from video recordings, questionnaire data and the robot's recordings of the behavioral data gave either partial or full support for our original hypotheses as formulated in Section 2. The physical embodied robot (K) in the gesture condition has been evaluated by the children as the interaction that they enjoyed the most. The drumming performance of the child–robot pair is the highest in the physical robot condition, and decreases in the virtual and disembodied robot conditions. Similarly, best turn-taking was achieved when they played with the physical robot; their coordination got worse in the virtual and disembodied robot conditions.

Results of questionnaire data analysis and behavioral data support the expectation that embodiment can play an important role in social interaction tasks. In particular, our child participants found the presence of a physical robot most enjoyable and believed it to be more pleasant to play with than a hidden robot or a virtual robot.

Overall, the questionnaire results show that children's opinions are not effected by the embodiment conditions, nor by the presence or absence of robot gestures. Nevertheless, it is interesting to note that the data indicates a trend in which the children generally appreciated more the interaction with the physically embodied robot than the other two conditions. Despite of that small result, significant interaction effects of embodiment and gesture conditions have been found in terms of enjoyment, intelligence and appearance.

In terms of enjoyment, there is a significant difference between the embodiment conditions when the robot made gestures during its drumming, where participants in the gesture condition enjoyed interacting with the embodied robot more than with the two other embodiment conditions. The result concerning intelligence (in which the perceived intelligence of the robot was less for the video and disembodied condition than for the embodied condition if gestures were used, while the opposite was true when no-gestures were performed) is also quite interesting, possibly highlighting the importance of physical embodiment for effective use of and the processing of non-verbal cues in social interactions. The result regarding appearance might suggest that the gestures used by the robot might be appropriate for the projected image, but not smooth enough for a robot sat in front of a child.

Moreover, the behavioral data of the children and the robot support that there is a significant difference between the embodiment conditions as we hypothesized. The presence of the physically embodied robot motivated the children positively, and helped to improve the turn-taking and drumming between the robot and the children significantly. When the robot made gestures whilst drumming, the differences between the drumming and turn-taking performances belonging to different embodiment conditions increased significantly. Gestures played a positive role, especially when the child played with the physically embodied robot in terms of turn-taking.

Also, there is a significant difference in the error rates of drumming and turn-taking between the first and the third games. Thus, as hypothesized, the children enjoyed the game more, and both the drumming and the turn-taking performance of the robot and the children improved as they played more. When the gestures were introduced, especially in the first game, there was a significant difference in the turn-taking errors — as the children played more, they got used to the robot and the gestures, and this difference decreased.

Note that we are aware of the limitations of our study. For example, the variances within the sample were quite large in comparison to the effect sizes. This is to be expected — the novelty of HRI scenarios, and considering that the sample consisted of children, would lead one to expect that idiosyncracies of the individual participant would impact both interactions and evaluations of these. Owing to this, a sample size such as the one presented in the paper (which is quite large compared to similar studies) is a reasonable and accepted way of controlling for such idiosyncracies in order to avoid a type II error ('Failing to reject a false null hypothesis') [47].

Several directions for future work can be envisaged. Considering that play is often a 'group activity', one may study teams of children playing together with the robot in order to investigate how they interact with each other as well as with the robot. Such 'social facilitation effects' in HRI have been found, for example, in our previous studies involving children playing games with a non-humanoid mobile robot [48]. Also, the experience of the human–humanoid drumming game could be compared with a human–human drumming game. Furthermore, adding visual feedback with the use of the robot's internal cameras may be helpful in enhancing the robot's interaction with the children. In that way the robot could adapt itself to the behavioral changes in the children to achieve a higher level of social interaction with them. Finally, different robot appearances (e.g., humanoid *versus* mechanical-looking or zoomorphic) and different robot behaviors could be varied systematically in studies comparing virtual and physical embodiment conditions.

On a general note, recently more emphasis has been given to the use of virtual social agents (e.g., virtual characters) in the context of interaction with humans [49–52]. Graphical characters in virtual environments (*cf*; computer games) can create scenarios far more complex than typical HRI scenarios, which explains their popularity in entertainment and educational applications, for example. However, interaction with virtual characters typically requires the use of dedicated interfaces, e.g., keyboard, mouse, Wii, etc., while interactions with robots, as described in this paper, do not necessarily require specialized input devices. As the results in this paper indicate, there is a clear benefit of using physical embodied characters, compared to virtual characters, as far as children's responses are concerned. Thus, despite limitations concerning the complexity of autonomous intelligent behavior and interactive capabilities of state-of-the-art robots, our research supports the need for physically embodied interaction in suitable scenarios. Note that in therapy applications the physical dimension of interaction can provide an additional strong incentive and therapeutic objective, and not unsurprisingly, more and more embodied robots have been used in special education, engaging children with special needs in meaningful interactions [53–55].

This large-scale study with an autonomous humanoid robot and children is one of the first studies in this domain of comparing physical and virtual robot embodiments. The achievements and findings here suggest implications for research in application areas involving robots and children. Our results indicate that the embodiment of the robot (virtual or physical) has a significant impact on the objective performance and children's subjective evaluation of the interaction.

Acknowledgements

This work was conducted within the EU Integrated Project RobotCub (Robotic Open-architecture Technology for Cognition, Understanding and Behaviours) funded by the EC through the E5 Unit (Cognition) of FP6-IST under Contract FP6-004370. We would like to thank our colleagues in the Adaptive Systems Research Group for their help in organizing the experiments: Scott Watson, Josh Wainer, Qiming Shen, Dr Michael L. Walters and Dr Joe Saunders. We would like to thank Dr René te Boekhorst for suggestions concerning the data analysis and Dr Frank Broz for proofreading the manuscript. Also, we are thankful to the schools who attended our event and the teachers who voluntarily helped us during the experiments.

References

1. K. Shinozaki, A. Iwatani and R. Nakatsu, Construction and evaluation of a robot dance system, in: *Proc. 12th Int. Conf. on Knowledge-based Intelligent Information and Engineering Systems*, Zagreb, pp. 682–691 (2008).
2. A. Tiago, C. Martinho, I. Leite and A. Paiva, iCat, the chess player: the influence of embodiment in the enjoyment of a game, in: *Proc. 7th Int. Conf. on Autonomous Agents and Multiagent Systems*, Estoril, pp. 1253–1256 (2008).
3. D. Golubovic, B. Li and H. Hu, A hybrid software platform for Sony AIBO robots, in: *Proc. RoboCup 2003: Robot Soccer World Cup VII*, Padua, pp. 478–486 (2003).
4. F. Michaud and S. Caron, Roball — an autonomous toy-rolling robot, in: *Proc. Workshop on Interactive Robot Entertainment*, Pittsburgh, PA, pp. 127–134 (2000).
5. A. Tapus and M. J. Matarić, Towards socially assistive robotics, *Int. J. Robotics Soc. Japan* **24**, 576–578 (2006).
6. K. Inoue, K. Wada and Y. Ito, Effective application of paro: seal type robots for disabled people according to ideas of occupational therapists, in: *Proc. ICCHP*, Linz, pp. 1321–1324 (2008).
7. B. Robins, K. Dautenhahn and J. Dubowsky, Does appearance matter in the interaction of children with autism with a humanoid robot? *Interact. Studies* **7**, 479–512 (2006).
8. K. Wada and T. Shibata, Robot therapy in a care house — its sociopsychological and physiological effects on the residents, in: *Proc. Int. Conf. on Robotics and Automation*, Orlando, FL, pp. 3966–3971 (2006).
9. I. Werry, K. Dautenhahn, B. Ogden and W. Harwin, Can social interaction skills be taught by a social agent? The role of a robotic mediator in autism therapy, in: *Proc. Int. Conf. on Cognitive Technology*, Coventry, pp. 57–74 (2001).
10. T. W. Bickmore and J. Cassell, Social dialogue with embodied conversational agents, in: *Natural, Intelligent and Effective Interaction with Multimodal Dialogue Systems*, J. v. Kuppevelt, L. Dybkjaer and N. Bernsen (Eds), pp. 23–54. Kluwer Academic, New York, NY (2005).
11. T. W. Bickmore and R. W. Picard, Establishing and maintaining long-term human computer relationships, *ACM Trans. Comp.-Hum. Interact.* **12**, 293–327 (2005).
12. K. M. Lee, N. Park and H. Song, Can a robot be perceived as a developing creature?: effects of artificial developments on social presence and social responses toward robots in human–robot interaction, Paper presented at: *Int. Communication Association Conf.*, New Orleans, LA (2004).
13. T. Fong, I. Nourbakhsh and K. Dautenhahn, A survey of socially interactive robots, *Robotics Autonomous Syst.* **42**, 143–166 (2003).
14. C. Lathan, J. Vice, M. Tracey, C. Plaisant, A. Druin, K. Edward and J. Montemayor, Therapeutic play with a storytelling robot, in: *Proc. Conf. on Human Factors in Computing Systems*, Seattle, WA, pp. 27–28 (2001).

15. J. Wainer, D. J. Feil-Seifer, D. A. Shell and M. J. Mataric, Embodiment and human–robot interaction: a taskbased perspective, in: *Proc. Int. Conf. on Human–Robot Interaction*, Jeju Island, pp. 872–877 (2007).
16. T. Quick, K. Dautenhahn, C. Nehaniv and G. Roberts, Understanding embodiment, system–environment coupling and the emergence of adaptive behaviour, in: *Intelligence Artificielle Située*, A. Drogoul and J.-A. Meyer (Eds), pp. 13–31. Hermes Science, Paris (1999).
17. K. Dautenhahn, B. Ogden and T. Quick, From embodied to socially embedded agents — implications for interaction-aware robots, *Cognit. Syst. Res.* **3**, 397–428 (2002).
18. R. L. Chrisley and T. Ziemke, Embodiment, in: *Encyclopedia of Cognitive Science*, L. Nadel (Ed.), pp. 1102–1108. Macmillan, London (2002).
19. K. M. Lee, Y. Jung, J. Kim and S. R. Kim, Are physically embodied social agents better than, disembodied social agents?: the effects of physical embodiment, tactile interaction, and people's loneliness in human–robot interaction, *Int. J. Hum.-Comput. Stud.* **64**, 962–973 (2006).
20. C. Bartneck, eMuu — an embodied emotional character for the ambient intelligent home, *PhD Thesis*, Eindhoven (2002).
21. A. Tapus and M. J. Matarić, Socially assistive robotic music therapist for maintaining attention of older adults with cognitive impairments, in: *Proc. AAAI Fall Symp. AI in Eldercare: New Solutions to Old Problem*, Washington, DC, pp. 297–298 (2008).
22. J. Wainer, D. J. Feil-Seifer, D. A. Shell and M. J. Mataric, The role of physical embodiment in human–robot interaction, in: *Proc. IEEE Int. Workshop on Robot and Human Interactive Communication*, Hatfield, pp. 117–122 (2006).
23. A. Powers, S. Kiesler, S. Fussell and C. Torrey, Comparing a computer agent with a humanoid robot, in: *Proc. ACM/IEEE Int. Conf. on Human–Robot Interaction*, Washington, DC, pp. 145–152 (2007).
24. C. Breazeal, *Designing Sociable Robots*. MIT Press, Cambridge, MA (2002).
25. C. Breazeal, Toward sociable robots, *Robotics Autonomous Syst.* **42**, 167–175 (2003).
26. C. Crick, M. Munz and B. Scassellati, Synchronization in social tasks: robotic drumming, in: *Proc. IEEE Int. Symp. on Robot and Human Interactive Communication*, Hatfield, pp. 97–102 (2006).
27. G. Weinberg and S. Driscoll, Robot–human interaction with an anthropomorphic percussionist, in: *Proc. Int. ACM Computer Human Interaction Conf.*, Montreal, pp. 1229–1232 (2006).
28. S. Kotosaka and S. Schaal, Synchronized robot drumming by neural oscillator, in: *Proc. Int. Symp. on Adaptive Motion of Animals and Machines*, Montreal, QC (2000).
29. H. Kose-Bagci, K. Dautenhahn, C. L. Nehaniv, Emergent dynamics of turn-taking interaction in drumming games with a humanoid robot, in: *Proc. IEEE Int. Symp. on Robot and Human Interactive Communication*, Munich, pp. 346–353 (2008).
30. H. Kose-Bagci, K. Dautenhahn, D. S. Syrdal and C. L. Nehaniv, Drum-mate: A Human-Humanoid Drumming Experience, in: *Proc. IEEE-RAS Int. Conf. on Humanoid Robots*, Pittsburgh, PA, pp. 242–247 (2007).
31. M. P. Blow, K. Dautenhahn, A. Appleby, C. Nehaniv and D. Lee, Perception of robot smiles and dimensions for human–robot interaction design, in: *Proc. IEEE Int. Symp. on Robot and Human Interactive Communication*, Hatfield, pp. 469–474 (2006).
32. P. Schermerhorn, M. Scheutz and C. R. Crowell, Robot social presence and gender: do females view robots differently than males?, in: *Proc. ACM/IEEE Int. Conf. on Human–Robot Interaction*, Amsterdam, pp. 263–270 (2008).
33. K. Dautenhahn, M. Walters, S. Woods, K. L. Koay, C. L. Nehaniv, A. Sisbot, R. Alami and T. Siméon, How may I serve you?: a robot companion approaching a seated person in a help-

ing context, in: *Proc. ACM SIGCHI/SIGART Conf. on Human–Robot Interaction*, Salt Lake City, UT, pp. 172–179 (2006).

34. D. S. Syrdal, K.-L. Koay, M. L. Walters and K. Dautenhahn, A personalised robot companion? The role of individual differences on spatial preferences in HRI scenarios, in: *Proc. IEEE Int. Symp. on Robot and Human Interactive Communication*, Jeju Island, pp. 26–29 (2007).
35. P. Schermerhorn, M. Scheutz and C. R. Crowell, Robot social presence and gender: do females view robots differently than males?, in: *Proc. ACM/IEEE Int. Conf. on Human–Robot Interaction*, Amsterdam, pp. 263–270 (2008).
36. FearNot! event, available from: http://www.herts.ac.uk/news-and-events/latest-news/Antibullying-software-and-robots.cfm, last accessed: 3 June 2009.
37. S. Sabanovic, M. P. Michalowski and R. Simmons, Robots in the wild: observing human–robot social interaction outside the lab, in: *Proc. AMC 2006*, Antalya, pp. 576–581 (2006).
38. T. Kanda, M. Shiomi, Z. Miyashita, H. Ishiguro and N. Hagita, An affective guide robot in a shopping mall, in: *Proc. ACM/IEEE Int. Conf. on Human–Robot Interaction*, San Diego, CA, pp. 173–180 (2009).
39. I. R. Nourbakhsh, C. Kunz and T. Willeke, The mobot museum robot installations: a five year experiment, in: *Proc. IEEE/RSJ Int. Conf. on Intelligent Robots and Systems*, Las Vegas, NV, pp. 3636–3641 (2003).
40. N. Dahlback, A. Jonsson and L. Ahrenberg, Wizard of Oz studies — why and how, in: *Proc. Int. Conf. on Intelligent User Interfaces*, Orlando, FL, pp. 193–200 (1993).
41. B. Robins, K. Dautenhahn and P. Dickerson, From isolation to communication: a case study evaluation of robot assisted play for children with autism with a minimally expressive humanoid robot, in: *Proc. Int. Conf. on Advances in Computer–Human Interactions*, Cancun, in press (2009).
42. B. Robins, K. Dautenhahn, R. te Boekhorst and C. L. Nehaniv, Behaviour delay and robot expressiveness in child–robot interactions: a user study on interaction kinesics, in: *Proc. ACM/IEEE Int. Conf. on Human–Robot Interaction*, Amsterdam, pp. 17–24 (2008).
43. G. Metta, P. Fitzpatrick and L. Natale, YARP: yet another robot platform, *Int. J. Adv. Robotics Syst. (Special Issue on Software Development and Integration in Robotics)* **3**, 43–48 (2006).
44. R. Bencina and P. Burk, Portaudio, available from: http://www.portaudio.com/trac/wiki/
45. H. Kose and H. L. Akin, Object recognition in robot football using a one dimensional image, in: *Proc. Turkish Symp. on Artificial Intelligence and Neural Networks*, Antalya, pp. 291–300 (2001).
46. J. C. McCroskey and T. A. McCain, The measurement of interpersonal attraction, *Speech Monogr.* **41**, 261–266 (1974).
47. J. Cohen, A power primer, *Psychol. Bull.* **112**, 155–159 (1992).
48. S. Woods, K. Dautenhahn and C. Kaouri, Is someone watching me? Consideration of social facilitation effects in human–robot interaction experiments, in: *Proc. IEEE Int. Symp. on Computational Intelligence in Robotics and Automation*, Espoo, pp. 53–60 (2005).
49. J. Cassell, T. Bickmore, H. Vilhjálmsson and H. Yan, More than just a pretty face: affordances of embodiment, in: *Proc. Int. Conf. on Intelligent User Interfaces*, New Orleans, LA, pp. 52–59 (2000).
50. J. Rickel, Intelligent virtual agents for education and training: opportunities and challenges, *Lecture Notes Comput. Sci.* **2190**, 15–22 (2001).
51. L. Chiattaro, I. Ieronutti and S. Rigutti, Supporting presentation techniques based on virtual humans in educational virtual worlds, in: *Proc. Int. Conf. on Cyberworlds*, Singapore, pp. 245–252 (2005).

52. M. Sapouna, D. Wolke, N. Vannini, S. Watson, S. Woods, W. Schneider, S. Enz, L. Hall, A. Paiva, E. Andre, K. Dautenhahn and R. Aylett, Virtual learning intervention to reduce bullying victimization in primary school: a controlled trial, *J. Child Psychol. Psychiatr.*, in press (2009).
53. A. Cook, K. Howery, J. Gu and M. Meng, Robot enhanced interaction and learning for children with profound physical disabilities, *Technol. Disabil.* **13**, 1–8 (2000).
54. B. Robins, P. Dickerson and K Dautenhahn, Robots as embodied beings — interactionally sensitive body movements in interactions among autistic children and a robot, in: *Proc. IEEE Int. Workshop on Robot and Human Interactive Communication*, Nashville, TN, pp. 54–59 (2005).
55. F. Michaud, T. Salter, A. Duquette, H. Mercier, H. Larouche and F. Larose, Assistive technologies and child–robot interaction, in: *Proc. AAAI Spring Symp. on Multidisciplinary Collaboration for Socially Assistive Robotics*, Stanford, CA (2007).

Appendix A

Descriptive Data Related to Questionnaire Data

Table A.1.
Descriptive data of the embodiment effect

	Embodiment condition	Mean	SE	95% Confidence interval	
				Lower bound	Upper bound
Enjoyment	*K*	4.457	0.090	4.278	4.637
	V	4.262	0.125	4.013	4.511
	D	4.376	0.097	4.182	4.570
Performance	*K*	3.667	0.139	3.389	3.944
	V	3.735	0.158	3.418	4.051
	D	3.735	0.142	3.451	4.019
Appearance	*K*	4.212	0.099	4.014	4.410
	V	4.020	0.118	3.784	4.256
	D	3.985	0.105	3.775	4.195
Social attraction	*K*	4.111	0.113	3.884	4.338
	V	4.009	0.139	3.730	4.287
	D	4.008	0.133	3.742	4.273
Intelligence	*K*	4.325	0.105	4.114	4.535
	V	4.275	0.115	4.045	4.505
	D	4.274	0.123	4.028	4.520
Involvement	*K*	4.455	0.087	4.280	4.629
	V	4.340	0.104	4.132	4.547
	D	4.330	0.107	4.117	4.543

Table A.2.
Descriptive data of the gesture effect

	Gesture condition	Mean	SE	95% Confidence interval	
				Lower bound	Upper bound
Enjoyment	gesture	4.422	0.120	4.181	4.662
	no-gesture	4.308	0.118	4.071	4.545
Performance	gesture	3.581	0.164	3.253	3.908
	no-gesture	3.843	0.164	3.516	4.171
Appearance	gesture	4.040	0.118	3.804	4.276
	no-gesture	4.104	0.118	3.868	4.340
Social attraction	gesture	4.135	0.162	3.811	4.460
	no-gesture	3.949	0.160	3.630	4.269
Intelligence	gesture	4.219	0.123	3.973	4.464
	no-gesture	4.364	0.121	4.122	4.605
Involvement	gesture	4.401	0.120	4.161	4.641
	no-gesture	4.348	0.118	4.112	4.585

Table A.3.
Descriptive data of the gender effect

	Gender	Mean	SE	95% Confidence interval	
				Lower bound	Upper bound
Enjoyment	male	4.359	0.134	4.091	4.627
	female	4.368	0.109	4.149	4.586
Performance	male	3.519	0.180	3.158	3.879
	female	3.846	0.150	3.546	4.146
Appearance	male	4.000	0.130	3.740	4.260
	female	4.123	0.108	3.906	4.339
Social attraction	male	3.968	0.181	3.607	4.329
	female	4.090	0.147	3.795	4.384
Intelligence	male	4.314	0.133	4.048	4.580
	female	4.415	0.109	4.197	4.632
Involvement	male	4 269	0.137	3995	4 543
	female	4.308	0.112	4.084	4.531

Table A.4.
Descriptive data of the interaction between embodiment and gesture

	Embodiment condition	Mean	SE	95% Confidence interval	
				Lower bound	Upper bound
Enjoyment[a]					
gesture	*K*	4.688	0.128	4.432	4.943
	V	4.281	0.178	3.927	4.636
	D	4.297	0.138	4.020	4.573
no-gesture	*K*	4.227	0.126	3.976	4.479
	V	4.242	0.175	3.893	4.592
	D	4.455	0.136	4.182	4.727
Performance					
gesture	*K*	3.621	0.196	3.229	4.014
	V	3.682	0.224	3.234	4.129
	D	3.439	0.201	3.038	3.841
no-gesture	*K*	3.712	0.196	3.320	4.105
	V	3.788	0.224	3.340	4.235
	D	4.030	0.201	3.629	4.432
Appearance[a]					
gesture	*K*	4.162	0.140	3.881	4.442
	V	4.182	0.167	3.848	4.515
	D	3.778	0.149	3.480	4.075
no-gesture	*K*	4.263	0.140	3.982	4.543
	V	3.859	0.167	3.525	4.192
	D	4.192	0.149	3.894	4.489
Social attraction					
gesture	*K*	4.328	0.162	4.005	4.651
	V	4.078	0.198	3.681	4.475
	D	4.000	0.189	3.622	4.378
no-gesture	*K*	3.894	0.159	3.576	4.212
	V	3.939	0.195	3.549	4.330
	D	4.015	0.186	3.643	4.388
Involvement					
gesture	*K*	4.500	0.124	4.251	4.749
	V	4.422	0.148	4.126	4.717
	D	4.281	0.152	3.978	4.585
no-gesture	*K*	4.409	0.123	4.164	4.654
	V	4.258	0.146	3.966	4.549
	D	4.379	0.150	4.080	4.678
Intelligence[a]					
gesture	*K*	4.438	0.150	4.138	4.737
	V	4.156	0.164	3.829	4.484
	D	4.063	0.176	3.712	4.413
no-gesture	*K*	4.212	0.148	3.917	4.507
	V	4.394	0.161	4.072	4.716
	D	4.485	0.173	4.139	4.830

[a] Significant interaction effect.

Table A.5.
Descriptive data of the interaction between embodiment and gender

	Embodiment condition	Mean	SE	95% Confidence interval	
				Lower bound	Upper bound
Enjoyment					
male	*K*	4.442	0.149	4.145	4.740
	V	4.269	0.197	3.876	4.663
	D	4.365	0.154	4.057	4.674
female	*K*	4.462	0.122	4.218	4.705
	V	4.256	0.161	3.935	4.578
	D	4.385	0.126	4.133	4.636
Performance					
male	*K*	3.352	0.211	2.930	3.774
	V	3.556	0.246	3.064	4.047
	D	3.648	0.229	3.190	4.106
female	*K*	3.885	0.176	3.533	4.236
	V	3.859	0.205	3.450	4.268
	D	3.795	0.191	3.414	4.176
Appearance					
male	*K*	4.198	0.155	3.887	4.508
	V	3.889	0.186	3.517	4.261
	D	3.914	0.169	3.576	4.251
female	*K*	4.222	0.129	3.964	4.480
	V	4.111	0.155	3.802	4.420
	D	4.034	0.141	3.753	4.315
Social attraction[a]					
male	*K*	4.212	0.184	3.844	4.579
	V	3.808	0.218	3.372	4.244
	D	3.885	0.209	3.467	4.302
female	*K*	4.038	0.150	3.739	4.338
	V	4.141	0.178	3.785	4.497
	D	4.090	0.171	3.749	4.431
Involvement					
male	*K*	4.462	0.138	4.185	4.738
	V	4.250	0.164	3.922	4.578
	D	4.231	0.168	3.895	4.566
female	*K*	4.449	0.113	4.223	4.674
	V	4.397	0.134	4.129	4.666
	D	4.397	0.137	4.123	4.671
Intelligence					
male	*K*	4.462	0.166	4.129	4.794
	V	4.231	0.183	3.865	4.597
	D	4.115	0.197	3.721	4.510
female	*K*	4.231	0.136	3.959	4.502
	V	4.308	0.150	4.009	4.606
	D	4.385	0.161	4.062	4.707

[a] Significant interaction effect.

Descriptive Data Related to Behavioral Data

Table A.6.
Descriptive data of the embodiment effect

	Embodiment condition	Mean	SE	95% Confidence interval	
				Lower bound	Upper bound
Errorsum[a]	*K*	0.623	0.063	0.497	0.750
	V	0.938	0.082	0.774	1.102
	D	0.899	0.095	0.710	1.088
Errorturn	*K*	0.741	0.016	0.709	0.772
	V	0.778	0.017	0.744	0.812
	D	0.780	0.017	0.746	0.814
Diffsum[a]	*K*	12.530	1.356	9.822	15.239
	V	18.379	1.622	15.139	21.619
	D	18.197	2.011	14.180	22.214
Maxofbeats[a]	*K*	5.500	0.261	4.978	6.022
	V	7.515	0.571	6.375	8.655
	D	6.576	0.367	5.843	7.308
Avgofbeats[a]	*K*	3.129	0.101	2.927	3.332
	V	3.932	0.219	3.494	4.369
	D	3.624	0.130	3.365	3.883

[a] Significant interaction effect.

Table A.7.
Descriptive data of the interaction between embodiment and gender

	Embodiment condition	Mean	SE	95% Confidence interval	
				Lower bound	Upper bound
Maxofbeats[a]					
male	*K*	5.630	0.408	4.814	6.445
	V	9.407	0.841	7.727	11.088
	D	7.815	0.542	6.732	8.898
female	*K*	5.410	0.340	4.732	6.089
	V	6.205	0.700	4.807	7.603
	D	5.718	0.451	4.817	6.619

[a] Significant interaction effect.

Table A.8.
Descriptive data of the interaction between embodiment and gesture

	Embodiment condition	Mean	SE	95% Confidence interval	
				Lower bound	Upper bound
Errorsum[a]					
gesture	*K*	0.642	0.090	0.463	0.822
	V	0.977	0.116	0.745	1.208
	D	1.139	0.134	0.872	1.406
no-gesture	*K*	0.604	0.090	0.425	0.784
	V	0.899	0.116	0.667	1.130
	D	0.659	0.134	0.392	0.926
Errorturn[a]					
gesture	*K*	0.705	0.023	0.659	0.750
	V	0.820	0.024	0.772	0.868
	D	0.820	0.024	0.773	0.868
no-gesture	*K*	0.777	0.023	0.732	0.822
	V	0.735	0.024	0.687	0.783
	D	0.740	0.024	0.692	0.788
Diffsum[a]					
gesture	*K*	11.636	1.917	7.806	15.467
	V	16.606	2.294	12.024	21.188
	D	21.485	2.843	15.805	27.165
no-gesture	*K*	13.424	1.917	9.594	17.254
	V	20.152	2.294	15.569	24.734
	D	14.909	2.843	9.229	20.589
Avgofbeats[a]					
gesture	*K*	3.211	0.143	2.925	3.497
	V	4.510	0.310	3.891	5.129
	D	3.897	0.183	3.531	4.264
no-gesture	*K*	3.048	0.143	2.761	3.334
	V	3.353	0.310	2.734	3.972
	D	3.351	0.183	2.985	3.717

[a] Significant interaction effect.

Table A.9.
Descriptive data of the game order effect

	Game order	Mean	SE	95% Confidence interval	
				Lower bound	Upper bound
Errorsum[a]	1	1.002	0.094	0.815	1.190
	2	0.821	0.079	0.663	0.978
	3	0.638	0.068	0.502	0.773
Errorturn[a]	1	0.790	0.014	0.761	0.818
	2	0.770	0.017	0.736	0.803
	3	0.738	0.020	0.698	0.778
Diffsum[a]	1	19.940	2.095	15.755	24.125
	2	17.745	1.774	14.201	21.289
	3	12.826	1.338	10.153	15.499
Maxofbeats[a]	1	6.056	0.295	5.466	6.645
	2	6.489	0.395	5.700	7.277
	3	7.548	0.525	6.499	8.598
Avgofbeats[a]	1	3.439	0.116	3.208	3.670
	2	3.491	0.163	3.165	3.817
	3	3.920	0.206	3.509	4.331
Non-zeroturns[a]	1	18.899	0.577	17.747	20.051
	2	19.644	0.524	18.597	20.690
	3	18.285	0.567	17.152	19.417

[a] Significant interaction effect.

Table A.10.
Descriptive data of the interaction between game order and gesture

	Game order	Mean	SE	95% Confidence interval	
				Lower bound	Upper bound
Errorsum[a]					
gesture	1	1.233	0.133	0.968	1.498
	2	0.851	0.111	0.629	1.073
	3	0.675	0.096	0.483	0.867
no-gesture	1	0.771	0.133	0.506	1.036
	2	0.790	0.111	0.568	1.013
	3	0.600	0.096	0.408	0.793

[a] Significant interaction effect.

Table A.11.
Descriptive data of the interaction between game order and gender

	Game order	Mean	SE	95% Confidence interval	
				Lower bound	Upper bound
Maxofbeats[a]					
male	1	6.778	0.453	5.872	7.684
	2	6.926	0.607	5.714	8.138
	3	9.148	0.808	7.535	10.762
female	1	5.333	0.377	4.580	6.087
	2	6.051	0.505	5.043	7.060
	3	5.949	0.672	4.606	7.291
Avgofbeats[a]					
male	1	3.681	0.178	3.326	4.036
	2	3.624	0.251	3.122	4.125
	3	4.454	0.316	3.823	5.086
female	1	3.197	0.148	2.901	3.492
	2	3.358	0.209	2.941	3.776
	3	3.386	0.263	2.861	3.912
Non-zeroturns[a]					
male	1	17.926	0.887	16.155	19.697
	2	19.852	0.805	18.243	21.461
	3	17.519	0.872	15.777	19.260
female	1	19.872	0.738	18.398	21.345
	2	19.436	0.670	18.097	20.775
	3	19.051	0.725	17.603	20.500

[a] Significant interaction effect.

About the Authors

Hatice Kose-Bagci received her MS and PhD degrees from the Computer Engineering Department, Bogazici University, Turkey, in 2000 and 2006, respectively. During her MSc and PhD studies, she worked in several research projects involving vision, localization and multi-agent planning in robot soccer, as well as working as a Teaching Assistant and Instructor in Bogazici University. She is currently a Research Fellow at the University of Hertfordshire, working in the EU sixth Framework Project RobotCub. Her current research focuses on gesture communication and imitation in child-sized humanoid robots. Her research interests include autonomous mobile robots, social robotics, interaction, communication and imitation in artificial systems and robotics.

Ester Ferrari is a Research Fellow in the School of Computer Science at the University of Hertfordshire. She is involved in the European IROMEC (Interactive Robotic Social Mediators as Companions) project that investigates the potential use of an interactive, autonomous robotic toy in therapy and education for children with special needs. She completed her PhD in Ergonomics at the Department of Psychology, University of Torino, Italy, in 2007. In the past few years she has conducted research in the areas of human–robot interaction, human-computer interaction, UCD design, usability, ergonomics and psychology.

Kerstin Dautenhahn is a Research Professor in the School of Computer Science at the University of Hertfordshire, where she coordinates the Adaptive Systems Research Group. She received her PhD degree from the Biological Cybernetics Department, University of Bielefeld, Bielefeld, Germany, in 1993, and subsequently worked at GMD (now Fraunhofer) in St Augustin, Germany, the VUB-AI Lab in Brussels, Belgium, and the University of Reading, UK, before she joined the University of Hertfordshire, in 2000. She has pioneered research in social robotics, human–robot interaction and assistive technology, and has published more than 200 research articles. She has edited several books and frequently gives invited keynote lectures at international meetings. She regularly organizes conferences and has been Principal Investigator of her research team in several European projects on developmental robotics (RobotCub), robot companions (Cogniron and LIREC), educational virtual environments (eCircus), and robotics and assistive technology (IROMEC, RoboSkin). She is Editor-in-Chief of the journal *Interaction Studies: Social Behavior and Communication in Biological and Artificial Systems*, as well as an Associate Editor of *Adaptive Behavior*, *International Journal of Social Robotics* and *IEEE Transactions on Autonomous Mental Development*. She is a Fellow of the Royal Society for the Encouragement of Arts, Manufactures and Commerce (RSA).

Dag Sverre Syrdal received his BS in Psychology from Queen's University Belfast, in 2001, and his MSc in Research Methods and Data Analysis in Psychology from the University of Hertfordshire, in 2002. After working as an ESL teacher in China, he joined the Adaptive Systems Research Group, in 2006 where his work has centered around the planning, execution and analysis of user studies within the Cogniron project. His research interests include human–robot interaction (HRI), with an emphasis on socially acceptable interactions in human-centered environments as well as the role of individual differences in HRI, visual attention and the psychology of religion.

Chrystopher L. Nehaniv (BSc, University of Michigan, 1987; PhD, University of California at Berkeley, 1992) is a Research Professor of Mathematical & Evolutionary Computer Sciences at the University of Hertfordshire, where he founded the Adaptive Systems Research Group together with Kerstin Dautenhahn, in 2000. His main interests include interactive systems, complex adaptive systems, artificial life, emergence of language and communication in embodied robots interacting with humans, as well as abstract algebra (theory of semigroups groups and sequential machines) and its applications (especially algebraic biology). He is also an Associate Editor of the journals *BioSystems* and *Interaction Studies*.

An Experimental Investigation of Interference Effects in Human-Humanoid Interaction Games*

Qiming Shen, Hatice Kose-Bagci, Joe Saunders, Kerstin Dautenhahn
Adaptive Systems Research Group, School of Computer Science
University of Hertfordshire, College Lane, Hatfield, AL10 9AB. United Kingdom
{Q.Shen, H.Kose-Bagci, J.1.Saunders, K.Dautenhahn}@herts.ac.uk

***Abstract*— Investigating how people respond to and relate to robots is a multifaceted scientific challenge. This paper reports on an experimental investigation concerning movement interference effects between a human and a robot. We compare results with that obtained by Oztop et al. [1], however, in our study we used a small child-sized robot (KASPAR) with an overall human-like appearance. The experiment was conducted with both child and adult participants who interacted with a small humanoid robot using arm waving behaviours. The experimental setup was designed to be less constrained than in [1] with an emphasis on playful interaction. The experimental results did not show evidence for interference effects. This might be due to a more game-like and less constrained experimental environment or to the specific features of the robot or both. In addition to measurements of the variance of the movements, we investigated a measure for behavioural synchrony between human and robot movements based on the concept of information distance. The results of information distance analysis indicated that most of the human participants were affected by the robot's behavioural rhythms. While our experiments did not show a movement interference effect, we found behavioural adaptation of participants' movement timing to the robot's movements. Thus, the measure of behavioural synchrony that we introduced appears useful for complementing other measures (such as variance) previously used in the literature.**

I. Introduction

As robots move ever closer to our daily lives Human-Robot Interaction (HRI) has become an increasingly important field of research [2]. Modelling human-human interaction is an important approach to HRI, and may provide inspiration to how the communicative and interaction dynamics as well as mechanisms can be realized in human-robot interactions. Human beings commonly interact with each other via actions and language. It is important for humans to understand the underlying meaning when people observe actions and hear speech from others. Many researchers suggest that mirror neurons play a critical role in action and language understanding [3, 4, 5, 6, 7].

Following the discovery of mirror neurons in the premotor cortex of macaque monkeys [8, 9], which discharge when the subject performs an action and when the subject observes a similar action made by another agent, a great deal of research concerning the nature of the mirror neuron system has been carried out [10]. One finding was that a similar mirror neuron system also exists in human brains [1, 10]. It has been suggested that the mirror neurons facilitate the imitation of observed actions, which demonstrates a matching between the perceived action and its execution [11, 12]. There have also been studies of 'interference effects' which are thought to occur as a result of the co-activation of conflicting populations of mirror neurons and are exhibited when a subject is observing and performing incongruent behaviours (illustrated in table 1). These effects have been found in human-human interactions, however, it is thought that they may also occur in human-robot interaction when the robot is more human-like [1, 13, 14, 15]. Recent research also found that interference effects were present when participants were told that a moving dot which they observed was generated by a human and absent when the phenomenon was described as computer generated [16]. Therefore, we may hypothesize that if the interference effects exist in interactions between humans and humanoid robots, it might suggest that humans may perceive such robots as possessing some "human-like" qualities instead of regarding them as simple mechanical machines. Such research may also provide hints at what type of robots may be acceptable as social interaction partners.

TABLE I
INTERFERENCE EFFECT ILLUSTRATION

Agent	Interference Absence		Interference Presence	
Human	Horizontal	Vertical	Horizontal	Vertical
Robot	Horizontal	Vertical	Vertical	Horizontal

In Oztop et al.'s work [1] they describe a human-robot and human-human interaction experiment in which they successfully found an interference effect in human-robot interaction using the mechanically looking, but humanoid robot called DB. Earlier work by Kilner et al. [13] did not find interference effects in human-robot interaction when a robotic arm was used. Thus, it appears from the previous literature that the appearance (and associated movements) of robots may have an impact on the interference effect.

* This work was conducted within the EU Integrated Project RobotCub ("Robotic Open-architecture Technology for Cognition, Understanding, and Behaviours"), funded by the EC through the E5 Unit (Cognition) of FP6-IST under Contract FP6-004370.

A starting point for our research was to expand this line of research further and conduct the experiments with a 'social robot[1]' with not only a humanoid shape but a human (child) -like overall appearance.

The main motivations underlying the research presented in this paper were to replicate the interference experiments with a social robot having a human-like appearance in a less constrained and more playful interaction scenario, to investigate whether children and adults would respond differently in such conditions, and finally, to study whether synchronisation of human and robot movements could be observed. The detailed research questions of this experiment are described in section II below.

II. Experimental Setup

In July 2008 an experiment similar to that described by Oztop et al. [1] was carried out, but using a less constrained experimental framework. It has been previously found that an interference effect exists in human-human interaction [1, 13], therefore in our experiment we only concentrated on human-robot interaction. In addition, this experiment introduced new variable factors such as the effect of music and a comparison of two different age groups of participants.

A. Research Questions

In this experiment, we investigated the following four research questions:

1. Can an interference effect be found in a playful human-robot interaction experiment using a 'social robot'?
2. Will the use of music affect the participants' behaviour in the interaction experiment?
3. Can we find significant differences between children and adults in terms of their behaviour in the interaction games?
4. Will the rhythm of human behaviour be affected by the rhythm of the robot's behaviour?

The word 'rhythm' in this paper means "a strong, regular repeated pattern of movement or sound" [17].

Our expectations were as follows: As explained in section I the literature suggests an effect of robot appearance on the interference effect. We thus expected that a robot with even more human-like appearance features (compared to DB used in [1]), would elicit a strong interference effect. However, the more playful and less constrained setup of the interaction experiment may influence the outcome. The playfulness of the interaction with the robot was introduced due to their appropriateness for child participants. We expected that music, which emphasizes the robot's movement rhythm would strengthen the interference effect. Since different levels of engagement of children versus adults interacting with a robot could be expected, we hypothesized to find different results for children and adults. Finally, we expected to find that participants would adapt the rhythm of their movements to the robot since previous research with a different version of the same robot has shown that children adapt the timing of their movements to the robot's movements [18]. Our measure of synchrony for human and robot movements in interaction used a previously introduced and experimentally verified method [22].

B. Synchrony Measurement

The method we used for identifying these similar and synchronous actions employed the idea of similarity using *information distance*, previously described by Crutchfield [19] and based on *information theory* [20]. Information distance was used here to capture the spatial and temporal relationships between events.

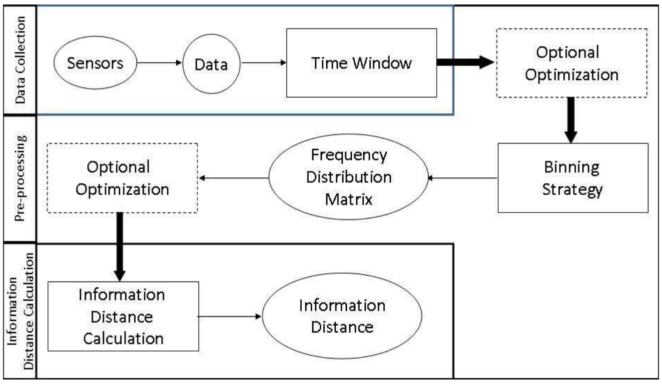

Fig. 1 The Similarity Method General Approach Flow Chart

The similarity identification method calculated the information distance between human and robot body part trajectories to yield an indication of their similarity. The numeric size of the information distance value gave an indication of similarity, the more similar the behaviors, the lower the value. Similarly, a higher value for information distance indicated less similar behaviors.

According to the general approach of this method (shown in Fig. 1), as a first step, the collected 3-D trajectory data of the participants and the robot movements was allocated into different data bins according to its value and the binning strategy. The binning strategy component was then used to extract data distribution features. These features were the critical source of information to conduct the information distance calculation. The calculation of information distance between two data columns, usually a pair of corresponding behavior components from the human and robot behavior respectively (for example, the x co-ordinates of the human forearm position and the x co-ordinates of the robot forearm position), is based on the information metric described by Crutchfield [19]. The information distance between two data columns X and Y is defined as the sum of two conditional entropies of these two columns [21]. It can be calculated using the following formula:

$$d(X,Y) = 2 * H(X,Y) - (H(X) + H(Y)) \quad [21]$$

This similarity identification model was verified using

[1] The term 'social robot' in the context of this paper refers to the humanoid robot KASPAR2 which has been designed by our research group with a number of human-like features and expressions (face, arms etc.) in order to facilitate human-robot interactions in 'social' contexts such as interaction games (as in this paper) or human-robot teaching. URL: http://KASPAR.feis.herts.ac.uk/

random data, artificial data, sine curve data and real human-robot interaction data. The validation results showed that the method was able to correctly identify similarity and synchronous behavior between a human and a robot, see more details in [22].

C. Experiment Design

The experiment described was conducted with both child and adult participants who interacted with a small humanoid robot. In total 14 children and 14 adults participated in the trials. However, following later video investigation, it was found that 4 child participants did not correctly follow the experimental instructions, which affected the data that was collected (e.g. one child tried to find out how fast the robot could move, rather than engaging in an interaction game). Therefore, the experimental data of these 4 children were excluded from the final data analysis. Note, all participants[2] were naive about the experiment.

The robot used in this experiment is called KASPAR2, developed by the Adaptive Systems Research Group at University of Hertfordshire. KASPAR2 is a child-sized humanoid robot with 18 DOF (degrees of freedom). It has 5 DOF in each arm, which enables it to perform some basic human-like waving behaviours. In this experiment, KASPAR2 only used its right arm (consistent with experiments in Oztop et al. [1]).

1) Waving Behaviours: Two basic waving behaviours were used in the experiment: vertical waving and horizontal waving. For both waving behaviours, the upper arm of a subject remained still and the subject used only the forearm, waving vertically or horizontally respectively. Therefore, the hand trajectory of the subject was curvilinear instead of linear, which was more natural and easy for both human and KASPAR2 to produce (note that in the Oztop et al.'s experiment [1] the trajectories were restricted to linear movements).

KASPAR2's waving behaviours were synchronized with a music track, which was the nursery rhyme: "Baa Baa Black Sheep". We chose a nursery rhyme because we expected that people may be more familiar with nursery rhymes and therefore find it easier to get involved in the music rhythm. In addition, many nursery rhymes have a slow and constant rhythm, which may allow better synchronization with KASPAR2's movements. The specified nursery music track had a duration of 30 seconds with a constant rhythm. The time interval between each beat in the music was 1.03 seconds and it took the robot 2.06 seconds to complete one single wave movement. That is, every single wave movement (for example, from left to right) of KASPAR2 took two beats and every complete back and forth wave movement (left to right then to left again) took four beats. During the whole experiment, KASPAR2 was waving at a constant speed. The transition between the with/without music conditions was conducted by simply switching on or off the computer speakers. With the music factor introduced, the participants were expected to synchronize more with the robot's behaviour when the music was on and to synchronize less when the music was off. Besides, music may make human-robot interaction more fun and more enjoyable.

2) Tracking System: A Polhemus Liberty magnetic motion tracking system was used to track the hand trajectories of both the human participant and of KASPAR2. Two magnetic sensors were attached on the waving hands of both human participants and KASPAR2 to collect data. The Liberty system returns the Cartesian coordinates of the sensors with respect to a fixed point (a large magnetic source).

TABLE II
EXPERIMENTAL SETUP COMPARISON

Experiment Setup Items		Oztop et al.'s Experiment	Current Experiment
Waving Behaviour	Direction	Top-Right to Bottom Left/Top-Left to Bottom Right	Vertical/Horizontal
	Frequency	0.5Hz	Not Specified
	Trajectory	Linear	Curvilinear
	Arm Used	Whole Arm	Forearm only
	Instructions Given	Detailed	General
Participants	Age	Adults	Adults/Children
	Distance to the Robot	2m	Around 1m
Agent		Robot/Human	Robot
Music		No Music	Music On/Off
Robot Platform			

3) Participant Instructions: During the experiment, the participants were asked to follow a few instructions. In order to create playful interaction, human participants were not specially trained to perform certain movements and many instructions given were very general instead of specifying every single detail:

1. Each participant was asked to stand facing KASPAR2 within a given distance (around one metre).
2. Each participant was shown the two basic waving movements described above and given a demo by the experimenter before starting the experiment.
3. Each participant was asked to only use their right arm in the experiment. However, the amplitude, speed and rhythm when the participant waved his or her arm was not restricted (from Oztop et al.'s [1] where the participants were explicitly instructed to be in phase with the other agent's movements).
4. Each participant was asked to concentrate on KASPAR2's waving arm when waving his or her arm.
5. Each participant was asked to interact with KASPAR2 for 8 trials.

These trials represented different experimental conditions according to 3 variables (2x2x2 within participant design, randomized order of the experimental conditions):

- arm waving direction (vertical/horizontal),

[2] The 10 children were all male and between 11 and 12 years old. The 14 adult participants (4 female, 10 male) were aged 18-52 (10 participants were between 23-26 years old). Thirteen adult participants were university students, one worked for a company.

- human-robot behaviour congruency (congruent/ incongruent) and
- music effect (with/without).

Each trial lasted around 30 seconds. Participants were informed when to start before each trial and when to stop after each trial.

The major differences in experimental setup between the experiment described in Oztop et al.'s work [1] and the current experiment are summarised in table 2.

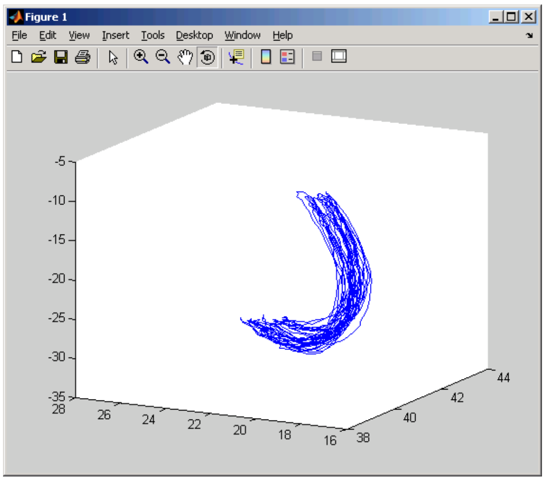

(a)

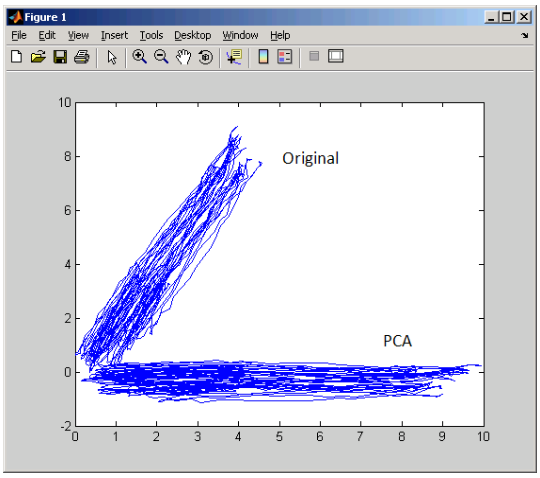

(b)

Fig. 2 (*a*) illustrates an example of human participant hand trajectory in 3-D space; (*b*) illustrates the mapping of the trajectory in figure *a* and the PCA of the mapping. The PCA is orthogonal to one of the axes.

III. Analysis of Results

A. Measurement Definition

The possible interference effects of human-robot interaction were measured by the variances in the waving movement, as in previous work e.g. by Oztop et al. [1]. In this experiment, the movement variances were defined as the variance orthogonal to a subject's main motion plane. For example, when a subject was waving horizontally, only the variances in the vertical direction (z-axis) were considered.

However, when a subject was waving vertically, it was more complex to locate the variances. This was because, in the experiment setup, the magnetic source was placed diagonally to the participants due to restrictions in the magnetic field generated by the Polhemus device. The range and position of the magnetic field also had to be limited to maintain the accuracy of measurement. Consequently, there was no axis (x, y or z) orthogonal to the subject's main motion plane in the vertical waving condition. An alternate approach applied was to take the mapped trajectory on the horizontal plane (x-y plane) and perform a PCA (Principal Components Analysis) to extract the desired axis. Usually, the first principle component (marked as the new x-axis, x') could be regarded as the mapping of the main motion plane on the horizontal plane. Therefore, the second principal component (marked as the new y-axis, y'), which was orthogonal to the x' axis, was the axis expected (Fig. 2). Through manual inspection, 94.8% of the vertical waving trajectories could use PCA to locate the axis. The axes of the rest of the trajectories were located manually.

In addition, the synchrony and similarity of the robot and participants' behaviours were also measured using an information distance approach [22], which was described in section II.

B. Interference Effect Analysis

A repeated-measures ANOVA was performed on the mean of the movement variances calculated across all trials for each condition (Table 3). Four fixed factors were involved in the ANOVA test: behaviour congruency, waving direction, presence of music and age group. The result showed that there was a significant effect in waving direction ($p < 0.05$) and age group ($p < 0.01$) (Fig.3).

However, there was no significant effect of congruency ($p > 0.1$) found in the experiment. The interaction effect between congruency and movement direction was not significant but very close ($p < 0.08$), which might potentially suggest that the congruent and incongruent behaviours had different impacts on the variability of the human movements in different directions (Fig. 3a).

The significant effect of waving direction was also found in Kilner et al.'s work [13] and Oztop et al.'s work [1], so our results validate their findings. Note, a possible explanation for the fact that we did not find support for the interference effect might be due to the different approaches used in locating the axis that the variance was calculated from.

The significant effect of age group suggested that the children and the adults behaved differently while interacting with the robot. The mean value of the variances in the children's behaviour was significantly higher than the adults' behaviour (Fig.3b). A possible explanation could be that the children adopted a stronger game-like attitude towards the task which lead to less constrained movements. Note, in the earlier work [1,13] higher variances have been interpreted as an indication for interference effects involving the mirror system. Our results did not show an interference effect but still higher variances in children's movements. Thus, future experiments need to investigate this finding further.

There was no significant effect overall in movement congruency. This may be due to the less constrained and more playful set up of the interaction experiment. The interference effect that might occur within a strict experimental setup might be overshadowed in a more relaxed and 'natural' human-robot interaction trial:

1. The type of the waving behaviour in our experiments was more natural (less linear).
2. The participants were not specifically trained to perform particular movements.
3. Only general instructions of how participants should wave their arms were given during the experiment.
4. There were no restrictions imposed on frequency or rhythm in participants' waving behaviours.

Thus, any of the factors mentioned above could have caused the interference effect to remain obscure in our experiments.

TABLE III

TESTS OF BETWEEN-SUBJECTS EFFECTS IN INTERFERENCE EFFECT ANALYSIS

Dependent Variable: Variance

Source	Type III Sum of Squares	df	Mean Square	F	Sig.
Corrected Model	54.489[a]	15	3.633	1.467	.122
Intercept	881.119	1	881.119	355.824	.000
Congruency	3.068	1	3.068	1.239	.267
Direction	16.783	1	16.783	6.777	.010
Music	.473	1	.473	.191	.662
Age	17.159	1	17.159	6.929	.009
Congruency * Direction	7.884	1	7.884	3.184	.076
Congruency * Music	4.338	1	4.338	1.752	.187
Congruency * Age	2.205	1	2.205	.890	.347
Direction * Music	.163	1	.163	.066	.798
Direction * Age	.283	1	.283	.114	.736
Music * Age	4.380E-5	1	4.380E-5	.000	.997
Congruency * Direction * Music	.073	1	.073	.029	.864
Congruency * Direction * Age	1.423	1	1.423	.575	.449
Congruency * Music * Age	2.042	1	2.042	.824	.365
Direction * Music * Age	.223	1	.223	.090	.764
Congruency * Direction * Music * Age	.159	1	.159	.064	.800
Error	435.824	176	2.476		
Total	1353.994	192			
Corrected Total	490.313	191			

a. R Squared = .111 (Adjusted R Squared = .035)

Besides, we did not find any significant effects for the music condition, which suggests that in our experiments music did not affect the variability of the human movements in human-robot interaction. Note, a possible explanation for this result could be that nursery rhymes may not have been suitable for either age group. However, we decided to chose one and the same music for both age groups, due to consistency purposes, and had assumed that both groups of participants may be familiar with such rhymes (e.g. via younger siblings or own children).

C. *Information Distance Analysis*

An ANOVA test was performed in the information distance analysis which was similar to the previous ANOVA test except the dependent variable was changed to information distance (Table. 4).

Significant effects were found in age group ($p < 0.01$), which validated the similar result in the variance interference effect analysis. Figure 5 shows that the mean value of information distance for children was much lower than the value for adults, suggesting the rhythm of waving in children's behaviour was more synchronized with the robot's rhythm than the adults' rhythm.

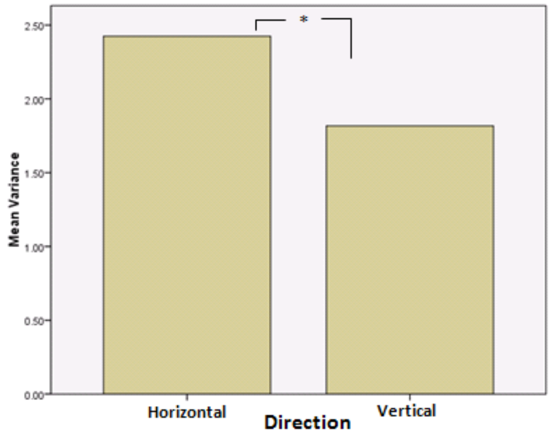

(a)

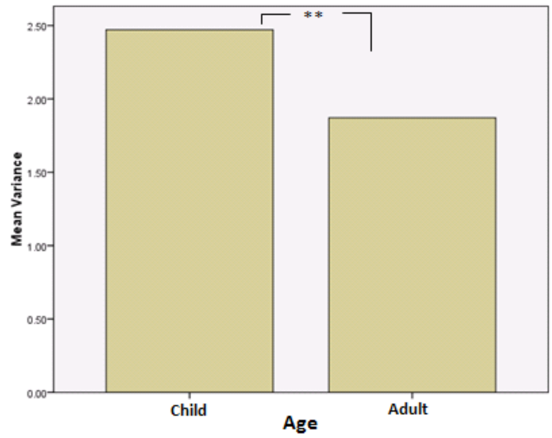

(b)

Fig. 3 figure *a* and *b* showed the effects of waving direction and age group. (a) The mean value of the variances that occurred in horizontal waving was much higher than the value of vertical waving. (b) The mean value of the variances that occurred in the behaviours of the children was much higher than the value in adults' behaviours. The significances of the ANOVA test are also shown in the figure (*: $p < 0.05$; **: $p < 0.01$).

A further statistical analysis of information distance values showed that the rhythm of waving behaviour of human participants was synchronized with the rhythm of the robot in over 81% of the trials (the information distance value of these trials were below 1.5, which was an empirical value indicating synchronization obtained in earlier research [22]). Note, that during the experiment, the participants were not instructed to wave with a particular rhythm or imitate the robot, instead, they were instructed to decide their behaviour rhythm by themselves. Therefore, the results show that the participants were affected by the robot's behaviour rhythm in the human-robot interaction experiments and adapted to it, which confirms previous results on timing adaptation in human-robot interaction experiments [18].

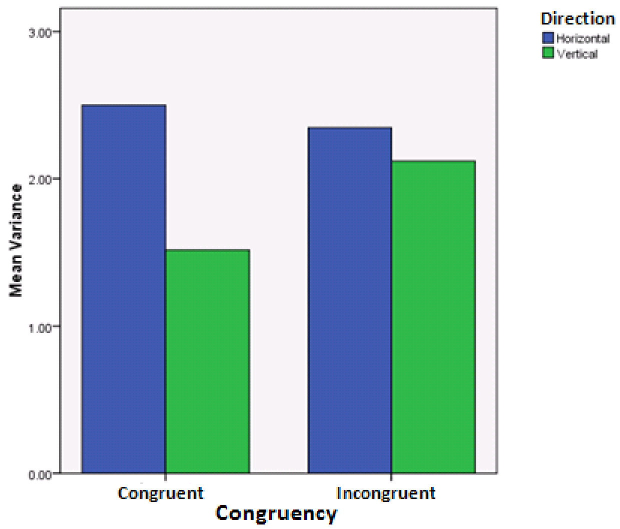

Fig. 4 The interaction effect between congruency and direction might potentially suggest that the congruent and incongruent behaviours had a different impact on the variability of the human movements in different directions

TABLE IV
TESTS OF BETWEEN-SUBJECTS EFFECTS IN INFORMATION DISTANCE ANALYSIS

Dependent Variable: InfoDist

Source	Type III Sum of Squares	df	Mean Square	F	Sig.
Corrected Model	3.643[a]	15	.243	1.534	.097
Intercept	280.766	1	280.766	1773.551	.000
Congruency	.059	1	.059	.371	.543
Direction	.000	1	.000	.001	.976
Music	.137	1	.137	.864	.354
Age	1.844	1	1.844	11.647	.001
Congruency * Direction	.014	1	.014	.090	.765
Congruency * Music	.046	1	.046	.291	.590
Congruency * Age	.000	1	.000	.002	.967
Direction * Music	.001	1	.001	.009	.925
Direction * Age	.248	1	.248	1.569	.212
Music * Age	.019	1	.019	.118	.732
Congruency * Direction * Music	.019	1	.019	.122	.728
Congruency * Direction * Age	1.176	1	1.176	7.431	.007
Congruency * Music * Age	.001	1	.001	.005	.943
Direction * Music * Age	.010	1	.010	.066	.798
Congruency * Direction * Music * Age	.055	1	.055	.344	.558
Error	27.862	176	.158		
Total	328.146	192			
Corrected Total	31.505	191			

a. R Squared = .116 (Adjusted R Squared = .040)

Note, Oztop et al. [1] relate their finding of the movement interference effect to the participants' perception of the robot as 'human'. In our experiments we did not find an interference effect, but we found behavioural adaptation of participants' movement timing to the robot. Thus, the measure of behavioural synchrony introduced above (section II) appears useful for complementing other measures (such as variance). This approach may offer a different route towards the multifaceted scientific challenge of understanding how people respond to and relate to robots.

There was no significant effect involving music in the information distance analysis. This may be because the rhythm of the music was the same as the behaviour rhythm of the robot. Thus, the facilitation effect of music could not be revealed even if it did exist.

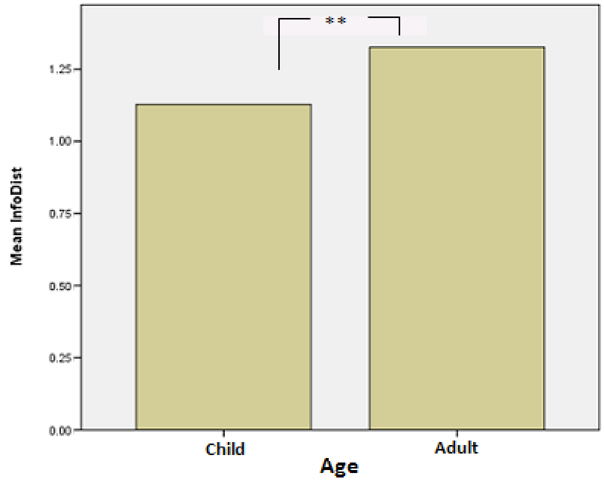

Fig. 5 This figure showed that the mean value of information distance of children was much lower than the value of adults. The significance of the ANOVA test is also shown in the figure (*: p <.0.05; **: $p < 0.01$).

IV. DISCUSSION AND FURTHER WORK

With respect to the research questions formulated in section II.A the results can be summarised as follows:

We did not find evidence for the movement interference effect in our experiments. This might be due to the less constrained and more playful experimental environment.

Alternatively, the specific robot used in the experiment could be an important factor. Recent research from neuro-imaging and neuro-psychological studies indicates that there are at least two routes of imitation: one is a goal-directed route and the other is a non-goal directed (the waving behaviours described in this paper can be regarded as non-goal directed behaviours as the participants were not informed of any particular goal during the interaction). The non-goal directed imitation appears to require from the imitatee greater reliance on effector selection (e.g. hand) and movement execution [7, 23]. Press et al.'s work [14, 15] may support this finding, which suggests that robotic stimuli have an impact on humans' mirror neuron systems if the robotic stimuli are similar to the human stimuli in visual properties. This implies that the limitations in robots can contribute to the absence of interference effects, which gives another possible explanation as to why there was no interference effect found in this experiment. There were some limitations in KASPAR2, which may affect the participants' concentration or behaviours during interactions:

1. The robot's servos were noisy.
2. The robot was in a sitting posture when the participants were standing, which caused differences in height between the robot and the participants. The participants were not instructed to sit on a seat because the seat would restrict the freedom of their behaviours.
3. There were temporary limitations in the robot neck servos. Therefore, it could not raise its head enough to face the participants.

An alternative explanation for the lack of the interference effect in the data is that the appearance of the robot is in the danger of falling into the 'uncanny valley' [24], which may be a factor in explaining whether robotic stimuli are effective or not. The "uncanny valley" is a theoretical idea that suggests that as robots become more human-like, they become less appealing to a real human. Only when true human-like features and movements appear does the "appeal" factor rise from the valley [24]. Although some robotic stimuli are very similar to human stimuli, if these stimuli fall into the uncanny valley and give humans a negative impression, then the mirror neuron system may not respond to them.

The experimental results indicated that the waving direction had significant impact on the human participants' behaviour, which validated similar results in Kilner et al. [13] and Oztop et al.'s work [1]. Our results also showed differences in movement variances between children and adults. In addition, the results of an information distance analysis indicated that most of the human participants were affected by the humanoid's behaviour rhythm, which may potentially suggest that the robot was regarded as an interaction 'partner'. We did not find any significant effect involving music. A possible explanation was that the rhythm of the robot's behaviours, which was the same as the music rhythm, are shadowed the effect of the music. Alternative, the choice of the music may have influenced the result.

Research into robot appearance suggests that an appropriate match between a robot's appearance and its social functionality can facilitate human acceptance and cooperation in interactions [25]. In this experiment, the servo noise and occasional shaky movements of KASPAR2 may have impaired its social functionality.

Moreover, some researchers found that children prefer interaction with a more machine-like robot over a more human-like robot [26, 27]. After the experiment, some participants, including children, reported that the rubber face of KASPAR2 looked scary. All of the feedback mentioned above indicated that KASPAR2 had very likely fallen into the uncanny valley, which may explain why KASPAR2 could not achieve responses from human participants' mirror neuron systems, although it looked more human-like (i.e. possessed more human-like appearance features) than e.g. the robot used in Oztop et al.'s work.

One may argue that the behavior rhythm could be affected by other simple rhythmic movements, e.g. caused by a pendulum or a moving dot on a screen instead of physical robots. That is, although the participants' behaviour rhythm was affected by the robot this may not necessarily mean that the participants treated the robot as a potential interaction partner. Our future work will try to validate this point by replicating the experiment using other visual stimuli instead of a robot. Further work may also change the rhythm of the music to further validate the impact of the music.

Acknowledgements

We would like to thank all the participants and especially the children of St. Matthew Academy, Blackheath, London for participating in the above study. We would like to thank Josh Wainer for assistance in the user study with the children.

References

[1] E. Oztop, D. W. Franklin, T. Chaminade, and G. Cheng (2005). "Human-humanoid interaction: is a humanoid robot perceived as a human?" in *International Journal of Humanoid Robotics* 2(4): 537-559.

[2] J. A. Adams and M. Skubic (2005), "Introduction to the special issue on human–robot interaction," in *IEEE Trans. Syst., Man, Cybern. C, Appl. Rev.*, 35(4): 433-437.

[3] G. Metta, G. Sandini, L. Natale, L. Craighero, L. Fadiga (2006), "Understanding mirror neurons: A bio-robotic approach," in *Interaction Studies* 7(2): 197-232.

[4] M. A. Arbib, (2002). "The mirror system, Imitation and the Evolution of Language." in *Imitation in animals and artifacts*, MIT Press.

[5] E. Oztop, M. Kawato and M. Arbib. (2006). "Mirror neurons and imitation: A computationally guided review." in *Neural Networks* 19: pp 254-271.

[6] G. Rizzolatti, M. A. Arbib, (1998). "Language within our grasp," in *Trends in Neurosciences*, 21(5): 188–194.

[7] Rizzolatti, G., Fogassi, L., & Gallese, V. (2001). "Neurophysiological mechanisms underlying the understanding and imitation of action," in Nature Reviews Neuroscience, 2: 661–670.

[8] G. Rizzolatti, L. Fadiga, V. Gallese, L. Fogassi (1996). "Premotor cortex and the recognition of motor actions." In *Brain Res Cogn. Brain Res.* 3: pp 131-141.

[9] V. Gallese, L. Fadiga, L. Fogassi and G. Rizzolatti (1996). "Action recognition in the premotor cortex," in *Brain* 119: 593-609.

[10] R. Hari, N. F., S. Avikainen, E. Kirveskari, S. Salenius, and G. Rizzolatti (1998). "Activation of human primary motor cortex during action observation: A neuromagnetic study," in *Proc. Nat. Acad. Sci.* USA 95: 15061-15065.

[11] S. Vogt, R. Thomaschke, (2007), "From visuo-motor interactions to imitation learning: Behavioural and brain imaging studies," in *Journal of Sports Sciences*, 25: 497–517

[12] E. Borenstein and E. Ruppin, (2005), "The evolution of imitation and mirror neurons in adaptive agents," in Cognitive Systems Research, 6(3): 229-242.

[13] J.M.Kilner, Y. Paulignan and S. J. Blakemore (2003), "An interference effect of observed biological movement on action," in *Current Biol* 13: 522-525.

[14] C. Press, G. Bird, R. Flach, & C. Heyes, (2005). "Robotic movement elicits automatic imitation," in *Brain Research: Cognitive Brain Research*, 25 (3): 632-640.

[15] C. Press, H. Gillmeister, C. Heyes. (2006). "Bottom-up, not top-down, modulation of imitation by human and robotic models," in *European Journal of Neuroscience*, 24(8): 2415-2419.

[16] E. Gowen, J. Stanley & R. C. Miall, (2008) "Movement interference in autism spectrum disorder," in *Neuropsychologia*, 46: 1060-1068

[17] www.askoxford.com (2008), "Compact Oxford English Dictionary," http://www.askoxford.com/concise_oed/rhythm?view=uk last accessed 28th Dec 2008

[18] B. Robins, K. Dautenhahn, R. te Boekhorst, C. L. Nehaniv (2008), "Behaviour Delay and Robot Expressiveness in Child-Robot Interactions: A User Study on Interaction Kinesics." In *Proc. ACM/IEEE 3rd International Conference on Human-Robot Interaction (HRI 2008).*

[19] J. P. Crutchfield (1990), “Information and its Metric,” in *Nonlinear Structures in Physical Systems – Pattern Formation, Chaos and Waves*, Springer Verlag, 1990, pp 119-130

[20] C. E. Shannon, “A mathematical theory of communication,” *Bell Systems Technical Journal,* vol. 27, pp. 379-423 and 623-656, 1948

[21] L. Olsson, C. L. Nehaniv, D. Polani (2006), “From Unknown Sensors and Actuators to Actions Grounded in Sensorimotor Perceptions,” in *Connection Science*, Vol. 18, Number 2, June 2006, pp. 121-144.

[22] Q. Shen, J. Saunders, H. Kose-Bagci, K. Dautenhahn (2008), “Acting and Interacting Like Me? A Method for Identifying Similarity and Synchronous Behavior between a Human and a Robot”, *Poster Presentation at IEEE IROS Workshop on "From motor to interaction learning in robots", September 26, 2008, Nice, France*.

[23] E. A. Franz, S. Ford, and S. Werner, (2007), "Brain and cognitive processes of imitation in bimanual situations: Making inferences about mirror neuron systems," in Brain Research, 1145: 138–149.

[24] M. Mori, (1970). “The Uncanny Valley,” in *Energy*, pp. 33-35.

[25] J. Goetz, S. Kiesler, A. Powers, (2003), “Matching robot appearance and behavior to tasks to improve human-robot cooperation,” in *Proceedings. ROMAN 2003*: 55-60

[26] B. Robins, K. Dautenhahn, R. te Boerkhorst, A. Billard (2004) "Robots as assistive technology - does appearance matter?" in *Proceedings, 13th IEEE International Workshop On Robot And Human Interactive Communication(ROMAN 2004)*: 277- 282

[27] S. Woods, K. Dautenhahn, J. Schulz, (2004), “The design space of robots: Investigating children's views,” in *Proceedings, 13th IEEE International Workshop On Robot And Human Interactive Communication(RO-MAN 2004*): 47-52.

Using Real-Time Recognition of Human-Robot Interaction Styles for Creating Adaptive Robot Behaviour in Robot-Assisted Play

Dorothée François, Kerstin Dautenhahn, Daniel Polani

Abstract— **This paper presents an application of the Cascaded Information Bottleneck Method for real-time recognition of Human-Robot Interaction styles in robot-assisted play. This method, that we have developed, is implemented here for an adaptive robot that can recognize and adapt to children's play styles in real time. The robot rewards well-balanced interaction styles and encourages children to engage in the interaction. The potential impact of such an adaptive robot in robot-assisted play for children with autism is evaluated through a study conducted with seven children with autism in a school. A statistical analysis of the results shows the positive impact of such an adaptive robot on the children's play styles and on their engagement in the interaction with the robot.**

I. Introduction

The work presented in this paper is part of the Aurora project, an ongoing long-term project investigating the potential use of robots to help children with autism overcome some of their impairments in communication, social interaction and imagination and fantasy[1]. Children with autism are able to play but the nature of their play may be described as restricted. Indeed, according to the American Psychiatric Association, "a lack of varied, spontaneous make-believe play is a defining feature of autism" [4]. Children with autism often play in a repetitive way, which can be linked to the children's preference for predictable environments. The advantage of enabling children with autism to interact with a robot is that robots enable simple and safe interaction by initially providing a relatively predictable environment for play. Progressively the complexity of the interaction can be increased.

Different possible obstacles have been identified that often prevent children with autism to actualize their potential for play. Among them are impairments in socioemotional intersubjectivity, impairment in joint attention and impairment in Theory of Mind [5]. These impairments negatively influence interaction in general and, more specifically, imply a lack of spontaneous and social reciprocity during play. Besides, the difficulty in perceiving the coherence of categories and concepts can be a reason why children with autism perceive an object in its parts and not as a whole, compare the weak central coherence theory [6; 7] for details. However, causes for impaired play are still not very well understood. These causes can vary for different children, depending also on the personality of the child and her past experience of play.

Yet play is an important vehicle for learning. Children can construct some understanding, i.e. active construction of meaning, through play. Besides, children usually enjoy playing (though this might not be the case in autism). Their pleasure and motivation seem to increase when they have the impression that they master a play situation [8]. Consequently, if we try to help children with autism master situations of play, they may have more fun playing which may contribute, even very modestly, to their quality of life. Play is also an important medium for self-expression [8].

Consequently, here we focus on facilitating play between children with autism and an autonomous robot, and particularly, we investigate the potential of a robot that can detect the children's play styles and adapt to them accordingly (such a robot is called an 'adaptive robot' in contrast to a 'reactive robot' which would only respond to current sensory input). Our goal is to encourage the children to engage in play and, when playing, to encourage 'well balanced' tactile interaction styles, i.e. neither too forceful nor too weak and within an intermediate frequency of interaction. We therefore address the following research questions:

- Does the adaptive robot, as described above, encourage or discourage the children from engaging in the interaction with the robot? Does their engagement change when interacting with a reactive robot?
- Does a child's play patterns differ when the robot is adaptive from when the robot is reactive? This question contains two subquestions as follows: i) Are the tactile strokes qualitatively different (ideally more gentle) when the child plays with an adaptive robot? ii) Is the frequency of the interaction differently (ideally better) balanced when the child plays with the adaptive robot?

In order to study these research questions, a hard technical challenge needs to be addressed, namely how to enable the robot to recognize in real time the tactile play styles of a child. This has been achieved by applying the Cascaded Information Bottleneck Method, a method that we developed and that is capable of extracting the temporal information of

Dorothée François, Kerstin Dautenhahn and Daniel Polani are with the Adaptive Systems Research Group, School of Computer Science, University of Hertfordshire, Hatfield AL10 9AB, UK (email: {d.francois, k.dautenhahn, d.polani}@herts.ac.uk).

The first author is supported by a research scholarship of the University of Hertfordshire. The work described in this paper was partially conducted within the EU Integrated Project RobotCub (Robotic Open-architecture Technology for Cognition, Understanding and Behaviours) and was partially funded by the European Commission through the E5 Unit (Cognition) of FP6-IST under contract FP6-004370. The authors would like to thank the teachers at the school, the children who took part in the experiments, and their parents. Thank you to René te Boekhorst for his useful input and comments on the statistical analysis.

[1]Autistic Spectrum Disorders can appear at various degrees and refer to different skills and abilities [1; 2]. Communication, social interaction and imagination and fantasy have been identified as the main impairments in autism, [3].

978-1-4244-2763-5/09/$25.00 ©2009 IEEE

a signal such as a time series of sensor data. We introduced it in a previous paper [9]. This method was developed as an extension of the well-known Information Bottleneck Method to the analysis of time series [10]. Section 3 briefly explains the method and provides details on its implementation for the recognition of human-robot interaction styles. We then report on trials conducted in a school with seven children with autism which evaluated the potential impact of such an adaptive robot on the children's play styles.

II. Related Work

Related work in robot-assisted play for children with autism has shown that when playing with a robot (in contrast to a stuffed animal), children with autism tend to show more behaviours that are typically impaired in autism (e.g. eye contact) [11]. Earlier comparisons between a mobile robot and a toy truck have shown more engaging behaviour towards the autonomous robot [12; 13]. Moreover other studies highlighted the potential role of the robot as a social mediator for children with autism [14; 15; 16]. Most studies were conducted in task oriented settings, e.g. involving imitation [16] or chasing games [15] with reactive (remotely controlled or autonomous) robots. Besides, the role of the experimenter in robot-assisted play has been investigated, firstly by Robins et al. [17] and more recently by François et al. [18].

The current paper focuses on investigating the role of an adaptive robot in robot-assisted play. We investigate whether an adaptive robot, i.e. a robot that could adapt to each child's play styles in real time would have a positive effect on the children's play styles and guide them progressively towards more well balanced interaction styles.

III. The Cascaded Information Bottleneck Method

A. Background: The Information Bottleneck Method

The Information Bottleneck Method [10] is a clustering method based on an information theoretic approach whose purpose is to extract the relevant information[2] in a signal $x \in \mathcal{X}$ that is, extract features of a random variable (r.v.) X that are relevant to the prediction of Y. This problem is modeled by the following Bayesian network with Markov condition: $\tilde{X} \longleftarrow X \longleftarrow Y$ where $\tilde{X}$ is the variable that extracts information about Y through X.

This popular method provides an alternative to 'rate distortion theory' techniques which constitute a standard approach to lossy source compression. In the Information Bottleneck method, the relevance is not addressed through an external distortion measure but directly through a variational principle implementing an information-theoretic formulation of sufficient statistics. The rationale is that the best trade-off between the compression of the signal and the preservation of the relevant information is the one that keeps a fixed amount of relevant information about the relevant signal Y while minimizing the number of bits from the accessible signal X, i.e. maximizing the compression. The optimal assignment can be found by minimizing the functional

$$\mathcal{L}[p(\tilde{x}|x)] = I(\tilde{X};X) - \beta I(\tilde{X};Y) \tag{1}$$

$I(X;Y)$ stands for the mutual information between X and Y. For β and the cardinality of $\tilde{X}$ fixed, an expression can be given which specifies implicitly the solution and leads to a fixed-point iteration. For the information bottleneck setting, the Kullback-Leibler divergence $D_{KL}(p(y|x)||p(y|\tilde{x}))$ replaces the distortion function from conventional rate-distortion theory.

The Agglomerative Information Bottleneck algorithm [19] makes the assumption that β tends to ∞ in the Lagrangian equation (Eq. 1). In this specific setting, the mutual information between $\tilde{X}$ and Y is maximized and a hard partition of the data into subsets is induced, each subset corresponding to a bottleneck state $\tilde{x}$: for a fixed cardinality of $\tilde{X}$ (i.e. a fixed number of subsets - also called states - in the bottleneck), each member of the input signal $x \in \mathcal{X}$ belongs to one and only one subset $\tilde{x} \in \tilde{\mathcal{X}}$ and $\tilde{x}$ is the subset for which $p(y|\tilde{x})$ has the smallest $D_{KL}(p(y|x)||p(y|\tilde{x}))$. The hard partition can be softened afterwards, with reverse annealing.

B. The Cascaded Information Bottleneck Method

1) The principle: Based on the Information Bottleneck Method, we have developed a novel time-filtering method particularly adapted for pattern recognition in time series. Let $x \in X$ be a time series input signal of length l, $x = [x_0, ..., x_{l-1}]$. We take k and $S \in \mathbb{N}$, with $l = k * S$, such that x can be divided into S disjoint sequences X_s, $s = 0, ..., (S-1)$, each of cardinality k, in the following way:

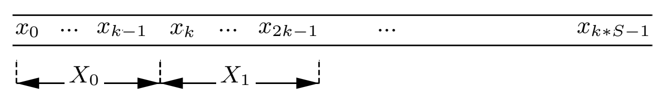

The Cascaded Information Bottleneck method relies on the principle that the relevant information can be progressively extracted from the time series with a cascade of successive bottlenecks sharing the same cardinality of bottleneck states but trained successively. The agglomerative information bottleneck algorithm is applied to each bottleneck successively, the first one being trained in the standard way while the next ones depend on the previous bottleneck states, as the following graph shows:

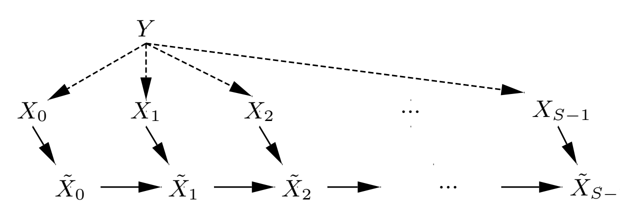

2) Extrapolation: The Cascaded Information Bottleneck Method progressively extracts the relevant information from an input sample $X = [X_0, ..., X_{S-1}]$ by a recall on the

[2] In this context, the relevant information is defined as the information that the (accessible) signal $x \in \mathcal{X}$ provides about another (typically not directly accessible) signal $y \in \mathcal{Y}$.

successive components (X_0 for the first step of the cascade, $(\tilde{X}_{s-1}, X_s)$ for step $s > 0$). Each bottleneck (we now discuss only $s > 0$, without loss of generality) is characterized by a probabilistic mapping $p(\tilde{x}_s|(\tilde{x}_{s-1}, x_s))$ which, for the present work using the agglomerative information bottleneck algorithm, is hard, i.e. above probability is 1 for exactly one value $\tilde{x}_s$ of $\tilde{X}_s$ and vanishes otherwise, i.e. it implements a hard mapping $(\tilde{x}_{s-1}, x_s) \mapsto \tilde{x}_s$ (note that the input $(\tilde{x}_{s-1}, x_s)$ corresponds to the input x of the original information bottleneck method).

During the information bottleneck training process, for each step of the cascade successively ($s > 0$), the mapping $(\tilde{x}_{s-1}, x_s) \mapsto \tilde{x}_s$ is built. If, however, at a step s in the cascade a pair $(\tilde{x}_{s-1}, x_s)$ never occurs during the training (we call this an *unseen pair*), the mapping $(\tilde{x}_{s-1}, x_s) \mapsto \tilde{x}_s$ will not be defined for the completed cascade. Upon processing of novel data, however, such a pair may be observed and in this case the cascade has no way to infer the following bottleneck state $\tilde{x}_s$, since there is no natural a priori correspondence of bottleneck states in successive bottlenecks.

For such cases, we therefore introduce an identification of successive bottleneck states which will provide us with a "default" continuation of a bottleneck state from step $s-1$ to step s in the case of unseen pairs. Let $\tilde{\mathcal{X}}_{s-1}$ and $\tilde{\mathcal{X}}_s$ be the set of bottleneck states $\tilde{x}_{s-1}$ and $\tilde{x}_s$, as well as $p(\tilde{x}_{s-1})$ and $p(\tilde{x}_s)$ their empirical probabilities. We consider one-to-one mappings r from $\tilde{\mathcal{X}}_{s-1}$ to $\tilde{\mathcal{X}}_s$ (which, for convenience, we call *permutations*). Each such permutation r provides an identification of successive bottleneck states. We define the informational cost of a permutation as

$$d_{(s-1,s)}(r) = -\sum_{\tilde{x}_{s-1} \in \tilde{\mathcal{X}}_{s-1}} p(\tilde{x}_{s-1}) \log \tilde{p}(\tilde{X}_s = r(\tilde{x}_{s-1})|\tilde{X}_{s-1} = \tilde{x}_{s-1}) \tag{2}$$

Note that $\tilde{p}(\tilde{X}_s = r(\tilde{x}_{s-1})|\tilde{X}_{s-1} = \tilde{x}_{s-1})$ is, for a given permutation r, the probability that the next state is $r(\tilde{x}_{s-1})$ knowing that the current state is $\tilde{x}_{s-1}$. The logarithm measures the unpredictability of the next state (i.e. the unpredictability of $\tilde{X}_s$ given $\tilde{x}_{s-1}$). If $\tilde{p}(\tilde{X}_s = r(\tilde{x}_{s-1})|\tilde{X}_{s-1} = \tilde{x}_{s-1}) = 0$ then, by convention, $d_{(s-1,s)}(r)$ is ∞.

To define a "default" continuation we now choose a permutation $R(s-1, s)$ that minimizes that unpredictability, weighted by the probability that the state $\tilde{x}_{s-1}$ actually occurs. Note that per construction of the bottleneck cascade, one never has $p(\tilde{x}_{s-1}) = 0$.

$$R(s-1, s) = \arg\min{}_r d_{(s-1,s)}(r) \tag{3}$$

$R(s-1, s)$ defines now a "default" path between $\tilde{X}_{s-1}$ and $\tilde{X}_s$, and thus provides an extrapolation of the succeeding bottleneck state in the case of an unseen pair.

3) Implementation: The Cascaded Information Bottleneck Method has been evaluated with two different criteria of interaction, namely the gentleness and the frequency of the interaction in [9]. The criterion gentleness contains two classes, namely 'gentle' and 'strong' which correspond respectively to non-forceful and forceful tactile interaction. The frequency of the interaction is categorised into four classes, defined by their typical periodicity of interaction: i) *very low* (S_0): the elapsed time between two tactile interactions is greater than 15 seconds; ii) *middle inferior* (S_1): the elapsed time between two tactile interactions is lower or equal to 15 seconds and greater than 5 seconds; iii) *middle superior* (S_2): the elapsed time between two tactile interactions is lower or equal to 5 seconds and greater than 1 second; iv) *very high* (S_3): the elapsed time between two tactile interactions is lower or equal to 1 second. S_1 and S_2 are considered here as well-balanced frequencies of interaction, while S_0 corresponds to a rare interaction and S_3 to a very intense interaction.

Two different cascades were built independently, one for each criterion of interaction. The gentleness corresponds to a short-term time scale event while the frequency corresponds to a mid-term time scale event (see Fig. 1 which provides the parameters for each cascade). The samples for the training of each cascade were generated during interactions with the Aibo ERS-7 in laboratory conditions within different runs. Each run contained one class exclusively, i.e. for the criterion gentleness, the samples generated within a same run contained only gentle or only strong styles of interaction (i.e. only gentle or only strong strokes were generated during a same run), and for the criterion frequency of the interaction, the samples generated within a same run contained only one type of frequency (i.e. S_0, S_1, S_2 or S_3 exclusively).

Criteria	Classes	Length of the input vector (window size), l	Length of the individual subsequences, k	Length of the cascade, S	Number of bottleneck states, m
Gentleness	2 classes: gentle/strong	50 (equivalent to 1.6 seconds)	2	25	4
Frequency	4 classes: S_0, S_1, S_2, S_3	472 (equivalent to 15.1 seconds)	2	236	6

Fig. 1. Parameters for each cascade of bottlenecks.

In both cases, a sliding window proceeds on the sensor data time series. For the criterion 'gentleness', the algorithm does not learn null samples (i.e. samples made of null events only). For the frequency of interaction, the system deals only with samples whose first component is not null.

The postprocessing relies on a 'winner takes all' principle: The selected (winner state) is defined by $\arg\max_{y \in Y} p(y|\tilde{x}_{S-1})$.

The method shows a sound recognition for both short-term and mid-term time scale events and involves only a very short delay for the recognition of short-term time scale events (0.17 seconds on average) [9]. Besides, the training process enables a structure to emerge over the cascade since the conditional entropy between the bottleneck states of two successive bottlenecks is globally decreasing over the cascade (Fig. 2). The Cascaded Information Bottleneck method is transparent and enables control over how much and what new information is taken at which step of the cascade. In particular, the extrapolation process enables to control the degrees of freedom of the system and prevent the cascade from over-learning.

In the next section, we present an application of the

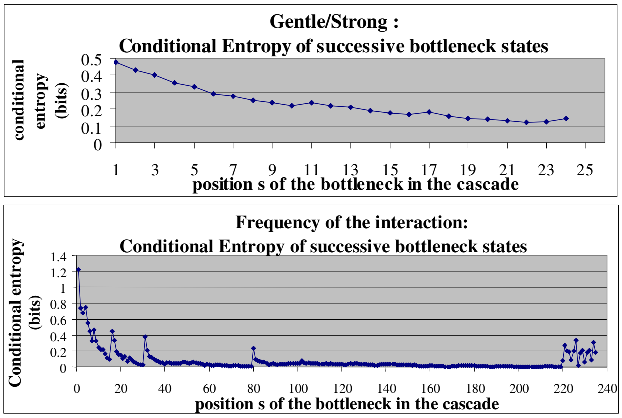

Fig. 2. Conditional entropy $H(\tilde{X}_{s+1}|\tilde{X}_s)$. $H(\tilde{X}_{s+1}|\tilde{X}_s)$ globally decreases over the cascade, pretty quickly, which suggests that a structure is progressively and rapidly emerging over the cascade: at the beginning of the cascade, a lot of new information is needed to deduce the next bottleneck state and then, when progressing in the cascade, less and less new information is needed. However, for the frequency of interaction, $H(\tilde{X}_{s+1}|\tilde{X}_s)$ has some small local peaks, both at the very beginning of the cascade and at the very end which suggest that at these steps s, the input data X_s may influence a bit more in the choice of next equivalent state $\tilde{X}_{s+1}$. Note that the ones at the end of the cascade may reflect the importance of the last steps for distinguishing the classes S_0 and S_1. In the rest of the study, the algorithm will extrapolate between step 5 and 24 (respectively 5 and 216) of the cascade for the gentleness (respectively frequency).

Cascaded Information Bottleneck method for Robot-Assisted Play whereby the method is used to enable a robot to recognize in real time human-robot interaction styles and adapt to them accordingly. In this application the criteria of interaction detected are the gentleness and the frequency of the interaction and the cascades are the ones described in the previous paragraph (and detailed in [9]).

IV. Application: A Real-Time Adaptive Robot for Robot-Assisted Play

A. The adaptive robot

In the context of this paper, a robot that is 'adaptive' can recognize interaction styles in real time and adapt to them appropriately. In other words, an adaptive robot reacts differently depending both on i) which sensor(s) is (are) activated (e.g. head sensor) and ii) the styles of interaction recognised. In contrast, by 'reactive' robot, we refer to a robot that can only react differently depending on which sensor is activated (e.g. head sensor or back sensor front), and which will not change its behaviour according to the interaction styles.

1) Reward of well balanced interaction styles: The adaptive mode relies on a reward basis for well-balanced interaction styles: the child should get a positive feedback from the robot when he/she plays in an appropriate style of interaction. The idea behind is that the child should always be encouraged and rewarded for every progress he/she made. With this approach, we hope to comfort the child in gaining self-confidence, enjoying himself/herself, and progressively acquiring a better understanding of the interactions he/she is involved in. It is hoped that the rewarding process can indirectly play the role of a trigger: the child wants to get the reward and therefore changes his/her behaviour until he/she actually gets it. Concretely, the robot should help regulate the interaction: if the child plays in a well-balanced interaction style, the robot reacts appropriately to the stimulation; on the contrary, if the interaction is e.g. too strong, the robot does not show any reaction. Moreover, the child should be encouraged to engage in the interaction if he/she is not engaged. Therefore, the robot should be both rewarding and engaging.

The reward is a physical reaction of the robot, which can be a gesture, a movement, a light or a sound. The concrete instantiation of these behaviours has been designed by immersion for each child beforehand during long-term studies with each child, whereby the experimenter tested different robot behaviours with each child in order to evaluate 1) whether the specific child liked it or not, 2) whether he/she conferred a specific meaning to the reaction and, particularly, whether the reaction had, in his/her view, a connotation of the robot being happy or sad.

We shall now detail the notion of reward: each time the child activates a sensor, the robot evaluates the interaction style in terms of gentleness and in terms of frequency and gives a reward, separately according to each criterion. If the interaction is gentle, then the robot shows a reaction to the child. The reaction depends on the sensor activated (there is a deterministic mapping between the sensors and the reactions of the robot for each child). If the stimulation takes place in a good overall frequency of interaction, i.e. a well-balanced frequency of interaction, then two LEDs turn on on the robot's face (which is sometimes interpreted by the children as the 'robot's eyes'). Note that a well-balanced frequency of interaction is a frequency not too low and not too high, represented in this study by the classes S_1 and S_2. Note, this model is generic and can be applied with different criteria of interactions. Fig. 3 presents the reward schema for the two criteria (gentleness and frequency) considered here.

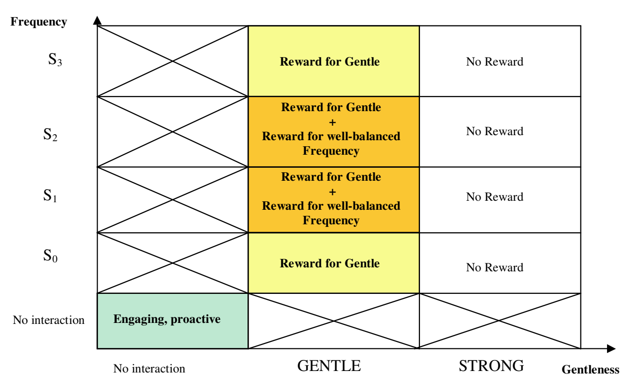

Fig. 3. Reward Schema for the two criteria of interaction.

2) Architecture for Decision-Making Based on Interaction Styles: The real-time recognition of the interaction styles uses the Cascaded Information Bottleneck Method. The small delay involved in the recognition process is modeled by a pause in the decision-making process, that is a small latency (600ms) during which the algorithm ignores the current interaction style. After the pause, the decision-making process looks at the successive classifications that are made by the Cascaded Information Bottleneck algorithm during a fixed short amount of time and counts the occurrences of strong behaviours recognised. If it exceeds a fixed predefined threshold then the final choice (i.e. the decision) is that the child's behaviour towards the robot is recognized as a 'strong interaction style' and the child will not get a reaction from the robot to his/her stimulation. If below threshold, then the decision 'gentle interaction style' is made and the child gets the reaction from the robot corresponding to the sensor activated. Besides, the robot updates the criterion frequency of interaction with a 1 second periodicity according to the Cascaded Information Bottleneck method (different threads for the gentleness and the frequency of interaction running in parallel). If, when the child strokes the robot gently, the current frequency of interaction is S_1 or S_2, then the child will get the additional reward of the two lights illuminating on the robot's face, while the robot also shows the specific reaction correlated to the gentle stimulation[3].

As for the evaluation of the child's disengagement, we consider that the child should be encouraged to play with the robot if he/she has not stroked the robot for a specific time that we define here as just above 15 seconds (more exactly, the length of the window size for classifying the frequency of the interaction which is 472×32 ms) which is reflected by a null input vector for the frequency of the interaction.

B. Trials

1) Participants: Seven children with autism participated in the experiments which took place in a school for moderate learning difficulties in UK. The children had had the chance to play with the robot during several months beforehand and were familiar with both the robot and the experimenter. The study was carried out with approval of the University of Hertfordshire Ethics Committee. The parents of all the children who took part in this study gave written consent, including permission to videotape the children.

2) Artifact: The robot was the Aibo ERS-7. It behaved autonomously and operated either in adaptive or reactive mode. In both cases the mapping between the sensors and the robot's reactions was the same except from the LEDs flashing for a good frequency of interaction, which was an additional feature for the adaptive robot, as well as wagging the tail when no interaction was detected. The behaviour mapping used for this specific study is detailed in Fig. 4.

3) Procedures and Measures:

Procedures: Each child participated in two sessions and the experiments involved one child at a time. Each session consisted of three successive steps[4] (also called games or runs), each step being defined by the mode of the robot– reactive (R) or adaptive (A)– which alternated between two successive steps.

[3]Note that this decision-making process really reflects the variety of the interaction styles considered here, the criterion 'gentle/strong' corresponding to a short-term time scale event and the criterion 'frequency of the interaction' corresponding to a mid-term time scale event.

[4]A session resulted in three steps also called games, which are, successively, step 1 (game 1), step 2 (game 2) and step 3 (game 3).

Sensor	Corresponding behaviour
Chin sensor	Emit "bark" sound while opening-closing the mouth
Head sensor	Turn head (Head tilt)
Back front sensor	- Wag the tail (used for Child E) - Walk forward, turn right, stand, turn left, walk backwards (used for the other children)
Back middle sensor	Turn head (Head pan)
Back rear sensor	Emit "drum" sound while wagging the tail

Fig. 4. Mapping between the external tactile sensors of the robot and its behaviours. For child E, the walking has been removed and replaced by the robot's wagging of the tail as part of the design by immersion.

As a result, a session was defined by its setting which was either A-R-A or R-A-R. Each child experimented with both settings (each during a different session, see Fig. 5).

Child	Setting 1	Setting 2
Child A	A-R-A	R-A-R
Child G	R-A-R	A-R-A
Child H	A-R-A	R-A-R
Child C	R-A-R	A-R-A
Child E	R-A-R	A-R-A
Child F	A-R-A	R-A-R
Child D	R-A-R	A-R-A

Fig. 5. Settings for the different children. Setting 1 corresponds to session 1 and setting 2 corresponds to session 2.

The robot's 'mode' was signaled to the child by a sticker with a specific geometrical form drawn on it (a triangle for adaptive and a circle for reactive mode); the sticker was put on the back of the robot at the beginning of each step. At each step, the child was told which game he/she was now playing, i.e. game 1 for step 1, game 2 for step 2 and game 3 for step 3. The child could see the experimenter putting the sticker on the back of the robot. The different stickers were used so that it was not too hard for the child to understand that the game was different (this procedure was considered to help the children cope with different experimental conditions). But the child had no information about the existence of adaptive and reactive modes; he/she could only possibly observe the difference in the reactions of the robot.

During each game, the child could freely interact with the robot. Before the start of each game, the experimenter:

1) paused the algorithm (for game 2 and 3),
2) congratulated the child and told him/her that now he/she would move on to game 2 (respectively 3),
3) put the corresponding sticker on,
4) sent the 'new robot's mode' through a wireless connection to the robot,
5) resumed the algorithm for the detection of play styles with the new robot's mode.

Each game lasted several minutes (depending on the children's specific needs and abilities); the minimum duration of each step was approximately 3 minutes. The experimenter did not touch the robot during the trials, except for putting on the sticker at the beginning of each step (sensor data were not collected at this stage), neither did she try to influence the child's behaviour in any way. The experimenter did not take part in the child-robot interactions in order not to interfere with the purpose of this study which had to focus on dyadic, uninterrupted interactions between the child and the robot, in order to test the potential of an adaptive robot to influence children's play styles.

Measures: The experiments were video recorded. The sensor data and the interaction styles detected with respect to the gentleness and the frequency of the interaction were recorded. These data were then analysed quantitatively off-line. For the criterion gentle/strong, we actually looked at the overall proportion of the sensor's activation and at the ratio of strong interaction styles. For the criterion 'frequency of the interaction', we took into account its evolution over time, which means here that we looked at the whole set of classifications, that is every 32 ms. We then used statistical techniques for non-parametric statistics.

C. Results

1) Statistical analysis of the engagement in the interaction and the gentleness of the strokes:

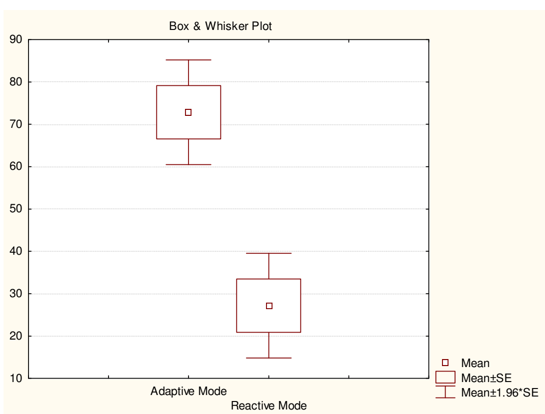

Fig. 6. Mean, Standard Error of the Mean (SE) and Confidence Intervals for the sensors' activation on the two modes. The x-axis represents the two modes; the y-axis represents the repartition in percentage of the sensors's activation.

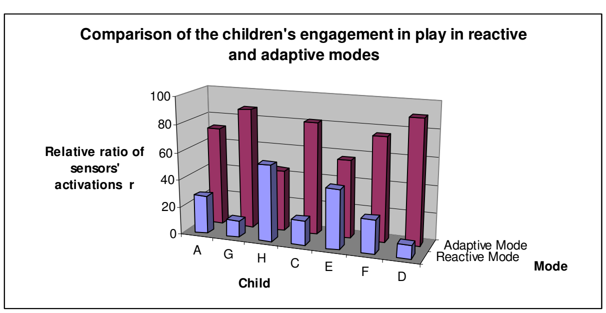

Fig. 7. Graph showing the relative engagement of the children in adaptive and reactive modes.

Engagement in the interaction: In order to study whether the adaptive robot may have a positive impact on the engagement of the children in play we do not consider the specificity of the strokes, i.e. whether they are gentle or strong. Instead, for each child we are interested in the total number of sensors' activations that we compare for adaptive and reactive robot modes.

For each child and for each mode, we count the total number of times the sensors were activated (each sensor[5] activated counts as one activation), namely, $N(Reactive)$, for the reactive mode, and $N(Adaptive)$, for the adaptive mode; for each child, we analyse the relative ratio of each mode[6], as follows:

$$r(Reactive) = \frac{N(Reactive)}{N(Reactive)+N(Adaptive)}$$

$$r(Adaptive) = \frac{N(Adaptive)}{N(Reactive)+N(Adaptive)}$$

The Wilcoxon test [20] is applied to the data from the seven children for the two following variables (Fig 7): $r(Adaptive)$, representing the adaptive mode, and $r(Reactive)$, representing the reactive mode. The test shows that there is a significant effect of the experimental conditions adaptive versus reactive (for $T = 1.000$, $p < 0.028$, with $N = 7$, Fig. 6). Thus, we can conclude that the children engage significantly more in the interaction during the adaptive mode..

Gentleness of the interaction: Here, we study the nature of the activation in terms of gentleness, i.e. whether an activation is gentle or strong. We therefore consider the percentage of strong strokes (also called strong activations) among the total number of sensor activations, per run and per child. For each child and for each mode, we take the average of this percentage over the runs from the two sessions[7] (Fig. 8).

Child	Average percentage of strong activations in the adaptive mode	Average percentage of strong activations in the reactive mode
Child A	20.52	71.97
Child G	2.08	12.50
Child H	5.56	9.09
Child C	3.53	11.75
Child E	15.23	15.79
Child F	17.51	67.74
Child D	60.58	33.33

Fig. 8. Table providing the average percentage of strong strokes in each mode for each child.

The Wilcoxon test is applied to the data from the seven children for the two following variables (Fig 8): the average of the percentage of strong strokes in the adaptive mode and the average of the percentage of strong strokes in the reactive mode. The test shows that there is no significant effect of the experimental conditions on the gentleness of the strokes ($N = 7$ and, for $T = 5.00$, on gets $p < 0.128$): there is no significant difference in the amplitude of the average

[5]Here we look at the activation of any of the four continuous external sensors, i.e. the head sensor and the three back sensors.

[6]Some children will naturally interact a lot with the robot, while others may stroke the robot only a few time during a session, thus we prefer to look at relative ratios.

[7]Here we consider the ratio of strong activations and investigate whether this ratio is inferior when the robot is in the adaptive mode, compared with when the robot is in the reactive mode.

percentage of strong strokes between adaptive and reactive modes. However, the proportion of cases where this average is smaller in the adaptive mode is 6 cases out of 7. The probability of obtaining such a deviation (6 or more cases out of 7) from a fifty-fifty ratio is 0.016 (two-tailed probability in the binomial test) which shows that, in the adaptive mode, the percentage of children who react less strongly in the adaptive mode deviates significantly from a fifty-fifty ratio.

2) Impact of the adaptive robot on the frequency of interaction: To analyse the impact of the adaptive robot on the frequency of interaction, we look at the four classes S_0, S_1, S_2, S_3 and how their occurrence varies in the adaptive and reactive modes.

We define R as the set of the three runs (steps) within a session for a specific child and $N_{S_i}(r)$ as the number of events from a class S_i for a specific run r. For each class S_i, each child, and each session, we define the relative ratio $\rho_{S_i}(r)$ for a given run r, defined as follows:

$$\rho_{S_i}(r) = \frac{N_{S_i}(r)}{\sum_{\tilde{r} \in R} N_{S_i}(\tilde{r})} \tag{4}$$

For each child, for each mode m (adaptive or reactive) and for each class S_i, the average relative ratio over the two sessions is called $Av_m(\rho_{S_i})$. For each child and for each mode m, the average relative ratio over the four classes is called $Av_m(\rho)$.

The Wilcoxon test is firstly applied to the two following variables: $Av_{Adaptive}(\rho)$ (representing the adaptive mode) and $Av_{Reactive}(\rho)$ (representing the reactive mode). The test shows that there is a significant effect of the experimental conditions (adaptive versus reactive) since for $T = 0$, one has $p < 0.018$, with $N = 7$. We can conclude that, in the adaptive mode, the interactions are significantly richer than in the reactive mode.

Secondly, the Wilcoxon test is applied for each class i separately, to the following variables: $Av_{Adaptive}(\rho_{S_i})$ (representing the adaptive mode) and $Av_{Reactive}(\rho_{S_i})$ (representing the reactive mode). For class S_0 (respectively class S_1) there is no significant difference between the two experimental conditions (adaptive versus reactive), since, for $T = 5.000$ (respectively $T = 4.000$), $p < 0.128$ (respectively $p < 0.173$) with $N = 7$. However, the proportion of cases where $Av_{Adaptive}(\rho_{S_0}) > Av_{Reactive}(\rho_{S_0})$ (respectively $Av_{Adaptive}(\rho_{S_1}) > Av_{Reactive}(\rho_{S_1})$) is 6 cases out of 7. The probability of obtaining such a deviation (6 or more cases out of 7) from a fifty-fifty ratio is 0.016 (two-tailed probability in the binomial test) which shows that the percentage of children for which there are more events related to S_0 (respectively S_1) in the adaptive mode than in the reactive mode deviates significantly from a fifty-fifty ratio. Concerning S_2 (respectively S_3) there is a significant effect of the experimental conditions Adaptive and Reactive since for $T = 1.000$ (respectively $T = 0.000$), $p < 0.028$ (respectively $p < 0.018$) with $N = 7$ (Fig. 9 and Fig. 10). Consequently, in the adaptive mode, there are significantly more events from class S_2 (respectively S_3) than in the reactive mode.

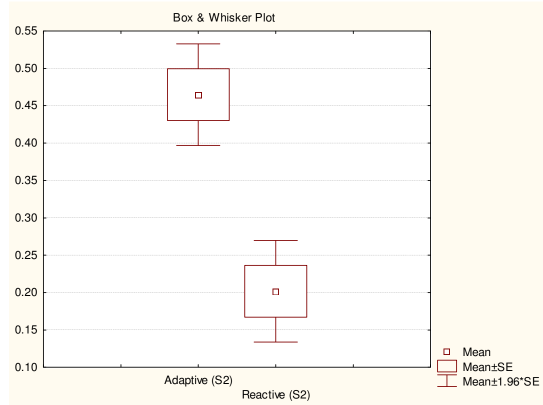

Fig. 9. Mean, Standard Error of the Mean (SE) and Confidence Intervals for S_2. The two variables are $Av_{Adaptive}(\rho_{S_2})$ and $Av_{Reactive}(\rho_{S_2})$.

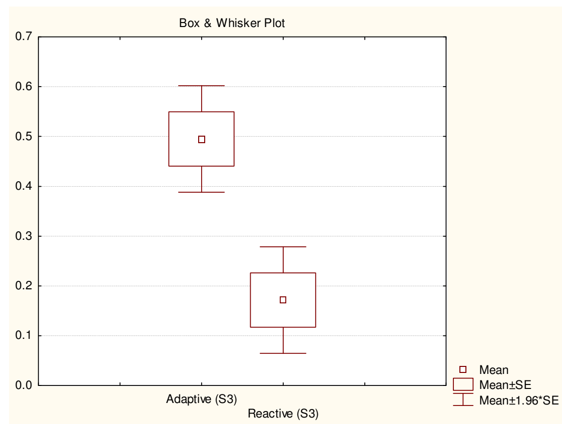

Fig. 10. Mean, Standard Error of the Mean (SE) and Confidence Intervals for S_3 The two variables are $Av_{Adaptive}(\rho_{S_3})$ and $Av_{Reactive}(\rho_{S_3})$..

V. Discussion

This study has shown that a robot that can adapt to child-robot interaction tactile styles can influence positively the children's play styles. Firstly, the children engaged significantly more in the interaction when the robot was adaptive and significantly more children played more gently with the robot in the adaptive mode. Besides, the interactions were significantly richer and higher frequencies including in particular a range of well balanced frequencies were significantly more present in the adaptive mode. The introduction of an adaptive robot in robot-assisted play for children with autism which is able to adapt in real time to children's interaction styles is a novel contribution. This is both a technical and methodological step forward in robot-assisted play. On the one hand, the development of a new computational method that enables robots to recognize in real time human-robot interaction styles is a step forward towards socially adaptive robots. On the other hand, the evaluation of the potential of an adaptive robot in robot-assisted play expands the role that robots could potentially play in the context of autism therapy.

The study conducted here involved only a few children and was short-term. Future work should expand this study with different children with autism. Note, in our study the children were familiar with both the robot and the experimenter and the robot's behaviour mapping had been tailored according to each child's needs and abilities which is very important in the context of autism. Future work should enable the same for the new children involved. Besides, future work should consider possible long-term effects of such an adaptive robot. In particular, future work could expand the model of adaptation by focusing on a larger grid of criteria for the interaction styles: while the child progresses, the robot could increase the range of criteria the child should meet to get a reward. In contrast, when the child encounters some difficulties, then the robot could simplify the range of criteria on which the reward for the child is based, so that the child could get a better understanding of the interactions happening. This progressive refinement in the adaptation process of the robot to the child's play styles could be linked, in some sense, to the notions of 'discrete development' and '(Alternate) Freezing and Freeing of Degrees of Freedom' which has been widely used in developmental robotics [21; 22; 23]. This technique, typically applied for a system learning motor skills, may be transposed to a social system, constituted here by the child and the robot: this social system is freezing some complexity in the interaction to learn more efficiently how to deal with interaction in general.

VI. Conclusion

This paper has presented an application of the previously introduced Cascaded Information Bottleneck Method for real-time recognition of Human-Robot interaction styles in the context of robot-assisted play for children with autism. Such an adaptive robot, which can detect the play styles of the children in real time and adapt to them accordingly, has been implemented and tested in the particular context of tactile interaction. The adaptation scheme rewards well balanced interaction styles and encourages the child to engage in the interaction if he/she is disengaged. The potential role of such an adaptive robot (compared to a reactive robot) in robot-assisted play has been evaluated with seven children with autism in a school and a statistical analysis showed that the adaptive mode influenced positively the play styles of the children in the following manner: 1) the children engaged significantly more in the interaction, 2) significantly more children played more gently with the robot, 3) the interactions were significantly richer and 4) the occurrence of higher frequencies, including a range of well balanced frequencies, was significantly increased. Future work should consider a wider study with different children with autism and investigate its long-term term impact. The Cascaded Information Bottleneck Method is generic and could potentially also be used in a variety of other applications in Robotics, Artificial Intelligence and Artificial Life. It is hoped that this work represents a step forward towards socially adaptive robots as well as robot-assisted play for children with autism.

References

[1] S. Powell, *Helping Children with Autism to Learn*. David Fulton Publish, 2000.

[2] R. Jordan, *Autistic Spectrum Disorders, An Introductory Handbook for Practitioners*. David Fulton publishers Ltd, 1999.

[3] "National autistic society," 2008, http://www.nas.org.uk/.

[4] A. P. Association, *Diagnostic and Statistical Manual of Mental Disorders DSM-IV, fourth Edition*. American Psychiatric Publishing, 1994.

[5] S. Baron-Cohen, *Mindblindness, an essay on Autism and Theory of Mind*. The MIT Press, new edition, 1997.

[6] U. Fritz, *Autism: Explaining the Enigma*. Oxford, Blackwell, 1989.

[7] K. Dautenhahn and I. Werry, "Towards interactive robots in autism therapy background, motivation and challenges," *Pragmatics and Cognition*, vol. 12(1), pp. 1–35, 2004.

[8] J. Boucher, "Editorial: interventions with children with autism - methods based on play," *Child language teaching and therapy*, vol. 15, pp. 1–5, 1999.

[9] D. François, D. Polani, and K. Dautenhahn, "Towards socially adaptive robots: A novel method for real time recognition of human-robot interaction styles," in *Proc. IEEE-RAS International Conference on Humanoid Robots (Humanoids08)*, 2008, pp. 353–359.

[10] N. Tishby, F. C. Pereira, and W. Bialek, "The information bottleneck method," in *Proc. of the 37-th annual Allerton Conference on Communication, Control and Computing*, 1999, pp. 368–377.

[11] C. M. Stanton, P. H. Kahn, R. L. Severson, J. H. Ruckert, and B. T. Gill, "Robotics animals might aid in the social development of children with autism," in *Proceedings of the 3rd International Conference on Human-Robot Interaction*, 2008.

[12] I. Werry, K. Dautenhahn, and W. Harwin, "Evaluating the response of children with autism to a robot," in *Proc. RESNA 2001*, 2001.

[13] I. Werry and K. Dautenhahn, *Human-Robot Interaction as a Model for Autism Therapy: An Experimental Study with Children with Autism.* In Modeling Biology: Structures, Behaviors, Evolution. Manfred Laubichler and Gerd B. Mller eds., Vienna Series in Theoretical Biology, MIT Press, 2007, pp. 283–299.

[14] H. Kozima, C. Nakagawa, and Y. Yasuda, "Interactive robots for communication-care: A case study in autism therapy," in *Proc. 14th IEEE Int. Workshop on Robot and Human Interactive Communication (RO-MAN)*, 2005, pp. 341–346.

[15] I. Werry, K. Dautenhahn, B. Ogden, and W. Harwin, "Can social interaction skills be taught by a social agent? the role of a robotic mediator in autism therapy," in *Proc. 4th Int. Conference on Cognitive Technology: Instruments of Minds (CT2001)*, 2001, pp. 57–74.

[16] B. Robins, K. Dautenhahn, and J. Dubowski, "Robots as isolators or mediators for children with autism? a cautionary tale," in *Proc. AISB'05 Symposium on Robot Companion Hard Problem and Open Challenges in Human-Robot Interaction*, 2005, pp. 82–88.

[17] B. Robins and K. Dautenhahn, "The role of the experimenter in hri research - a case study evaluation of children with autism interacting with a robotic toy," in *Proc. 15th IEEE Int. Symposium on Robot and Human Interactive Communication (RO-MAN 06)*, 2006, pp. 646–651.

[18] D. François, S. Powell, and K. Dautenhahn, "A long-term study of children with autism playing with a robotic pet: Taking inspirations from non-directive play therapy to encourage children's proactivity and initiative taking," *To appear in Interaction Studies. Special Issue: 'Robots in the Wild: Exploring Human-Robot Interactions in Naturalistic Environments'*, 2009.

[19] N. Slonim and N. Tishby, "Agglomerative information bottleneck," in *Proc. of Neural Information Processing Systems (NIPS 99)*, 1999, pp. 617–623.

[20] S. Siegel and N. J. Castellan, *Nonparametric Statistics for the behavioral sciences*. McGraw-Hill International Editions, 1988.

[21] L. Berthouze and M. Lungarella, "Motor skills acquisition under environmental perturbations; on the necessity of alternate freezing and freeing of degrees of freedom," *Adaptive Behaviour*, vol. 12(1), pp. 47 – 64, 2004.

[22] M. Lungarella and L. Berthouze, "Adaptativity via alternate freeing and freezing of degrees of freedom," in *Proc. 9th Int. Conf. on Neural Information Processing*, 2002, pp. 482–487.

[23] G. Gómez, M. Lungarella, P. E. Hotz, K. Matsushita, and R. Pfeifer, "Simulating development in a real robot: on the concurrent increase of sensory, motor, and neural complexity," in *Proc. 4th Int. Workshop on Epigenetic Robotics: Modeling Cognitive Development in Robotic Systems*, 2004, pp. 119–122.

2008 8th IEEE-RAS International Conference on Humanoid Robots
December 1 ~ 3, 2008 / Daejeon, Korea

Towards Socially Adaptive Robots: A Novel Method for Real Time Recognition of Human-Robot Interaction Styles

Dorothée François [1], Daniel Polani [1], Kerstin Dautenhahn [1]

[1] *Adaptive Systems Research Group*
Department of Computer Science
University of Hertfordshire
Hatfield AL10 9AB, UK
{d.francois, d.polani, k.dautenhahn}@herts.ac.uk

***Abstract*—Automatically detecting different styles of play in human-robot interaction is a key challenge towards adaptive robots, i.e. robots that are able to regulate the interactions and adapt to different interaction styles of the robot users. In this paper we present a novel algorithm for pattern recognition in human-robot interaction, the Cascaded Information Bottleneck Method. We apply it to real-time autonomous recognition of human-robot interaction styles. This method uses an information theoretic approach and enables to progressively extract relevant information from time series. It relies on a cascade of bottlenecks, the bottlenecks being trained one after the other according to the existing Agglomerative Information Bottleneck Algorithm. We show that a structure for the bottleneck states along the cascade emerges and we introduce a measure to extrapolate unseen data. We apply this method to real-time recognition of Human-Robot Interaction Styles by a robot in a detailed case study. The algorithm has been implemented for real interactions between humans and a real robot. We demonstrate that the algorithm, which is designed to operate real time, is capable of classifying interaction styles, with a good accuracy and a very acceptable delay. Our future work will evaluate this method in scenarios on robot-assisted therapy for children with autism.**

***Index Terms*—Socially interactive robots, socially adaptive robots, pattern recognition, human-robot interaction, robot-assisted play**

I. Introduction

This study is part of the Aurora project [1], an ongoing long-term project investigating the potential use of robots as a therapeutic toy for children with autism. One main stream of this project focuses on developing methods enabling the robot to analyze in real time the interaction styles and adapt its own behaviour appropriately with respect to a child's specific needs and abilities[1].

This paper presents a novel method for time series analysis, the Cascaded Information Bottleneck Method, which we apply to the real-time recognition of human-robot interaction styles. This method, which enables time-filtering, is based on the concept of Information as introduced by Shannon [2] and builds upon from the "Information Bottleneck Method" developed by Tishby et al. in [3].

Importantly, this work goes beyond prior work that either classified and characterized interactions off-line, i.e. after the interactions had taken place, or relied on explicit criteria tuned by hand (vs. automated training phase of the recognition algorithm). It also goes beyond previous work of the authors which enabled real-time recognition of interaction styles with respect to one criterion, the gentleness, using a different method, based on self-organizing maps [4]. The Cascaded Information Bottleneck Method is entirely generic for applications with socially interactive robots; in particular, it can be applied to humanoid-human interaction.

The remainder of the paper is structured as follows. Section II introduces related work. Section III summarizes some background on the Information Bottleneck Method developed by Tishby et al. in [3]. Section IV presents the Cascaded Information Bottleneck Method. The application to the recognition of Human-Robot Interaction Styles is explained in the two following sections, with details on the implementation and description of the trials in Section V and presentation of the results in Section VI. Section VII discusses the results and future work. Conclusion closes the paper (Section VIII).

II. Related Work

The role of tactile human-robot interaction in educational and therapeutic applications has been well highlighted by long-term studies with the seal robot Paro which have proven that specific everyday life situations exist in which human-robot interaction can have a positive effect on the well-being of human beings [5] and even play a role in a therapeutic context of cognitive and physical rehabilitation [6]. The Huggable robot, a teddy-bear like robot, equipped with a full body sense of touch, has proven to be a promising support to investigate the quantitative characterisation of the social affective content of touch [7]. Offline characterisation of interaction styles in general, moreover, has been investigated recently with diverse approaches. In [8], Scassellati focused on providing quantitative and objective measurements to assist in the diagnosis of autism. Measurements refer to the position in the room, vocal

[1] We consider the child's abilities as they are expressed through interaction with the robot, resulting in different play styles.

978-1-4244-2822-9/08/$25.00 ©2008 IEEE

prosody and gaze pattern – whose characterisation relies on linear discriminant analysis which is a clustering technique used for linearly separable data. Kanda et al. conducted a study [9] that highlighted the feasibility to link quantitative robot's and human's data characterizing body movements with a subjective evaluation made by the participant. Later, in [10] Salter et al. showed the possibility, in the context of child-robot interaction, to reflect some traits of personality of the children with an offline clustering technique based on the empirical probability distribution of the activation of the sensors.

Concerning real-time classification of interaction styles, in [11], Salter et al. have presented a real-time simple recognition algorithm for four interaction styles ('alone', 'interacting', 'carrying' and 'spinning') using the robotic platform Roball. The algorithm is based on a decision tree whose conditions are set up manually, by visual inspection of sensor data. In [12], Derakhshan et al. present an interesting real-time classification algorithm of interaction styles for children playing on an adaptive playground that is made of tiles equipped with sensors. The algorithm relies on a multi-agent system approach of BDI (Belief—Desire—Intention) in combination with neural networks using supervised learning. It shall be further noted that in the slightly different context of gesture recognition, Hidden Markov Models (HMMs) have been largely used for real-time recognition [13]. An HMM is defined by its number of hidden states and the two following probability matrices: the transition matrix, describing the conditional probabilities, given the state S at time step t, to be in the state S' at time $t+1$, and the emission matrix, defining the conditional probability of emitting a signal O, given the state S. Those matrices are static, i.e. for a given HMM, those values are fixed in time. Classifying an observation with HMMs consists in finding, among all the different HMMs[2] the one which has the highest probability of emitting this observation [14].

III. Background: The Information Bottleneck Method

The Information Bottleneck Method [3] is a clustering method based on an information theoretic approach [2] whose purpose is to extract the relevant information[3] in a signal $x \in \mathcal{X}$ that is, extract features of a random variable (r.v.) X that are relevant to the prediction of Y. This problem is modeled by the following Bayesian network with Markov condition: $\tilde{X} \longleftarrow X \longleftarrow Y$ where $\tilde{X}$ is the variable that extracts information about Y through X.

This method provides an alternative to 'rate distortion theory' techniques which constitute a standard analysis of lossy source compression. In the Information Bottleneck method, the relevance is not addressed through distortion but directly through a new variational principle. The rationale is that the best trade-off between the compression of the signal and the preservation of the relevant information is the one that keeps a fixed amount of relevant information about the relevant signal Y while minimizing the number of bits from the original signal X, i.e. maximizing the compression. The optimal assignment can be found by minimizing the functional

$$\mathcal{L}[p(\tilde{x}|x)] = I(\tilde{X};X) - \beta I(\tilde{X};Y) \qquad (1)$$

I(X;Y) stands for the mutual information[4] between X and Y. For β and the cardinal of $\tilde{X}$ fixed, an expression can be given which specifies implicitly the solution and leads to a fixed point iteration. β can be considered as the inverse of the temperature. This method uses a stochastic clustering top-down approach. The notion of stochastic refers to the fact that the clustering is soft and that the input data are mapped to the different elements of X with a particular probability. For that information bottleneck setting, the Kullback-Leibler divergence $D_{KL}[p(y|x)|p(y|\tilde{x})]$ replaces the distortion function.

The Agglomerative Information Bottleneck algorithm [17] makes the assumption that β is ∞ in the Lagrangian equation (1). It maximizes the mutual information between $\tilde{X}$ and Y and induces a hard partition of the data : for a fixed cardinal of $\tilde{X}$ (i.e. a fixed number of subsets - also called states - in the bottleneck), each member of the input signal $x \in \mathcal{X}$ belongs to one and only one subset $\tilde{x} \in \tilde{\mathcal{X}}$ and $\tilde{x}$ is the subset (the state) for which $p(y|\tilde{x})$ has the smallest $D_{KL}[p(y|x)||p(y|\tilde{x})]$. The hard partition can be soften afterwards, with reverse annealing. The pseudo-code of the algorithm can be found in [17].

IV. The Cascaded Information Bottleneck Method

A. The principle

Based on the Information Bottleneck Method, we have developed a novel time-filtering method particularly adapted for pattern recognition in time series. Let x $\in X$ be the time series input signal of length l, $x = [x_0, ..., x_{l-1}]$. We take k and $S \in \mathbb{N}$, with $l = k * S$, such that x can be divided into S disjoined parts of cardinality k, X_s, $s = 0, ..., (S-1)$ in the following way:

$$\underbrace{x_0 \ \cdots \ x_{k-1}}_{X_0} \ \underbrace{x_k \ \cdots \ x_{2k-1}}_{X_1} \ \cdots \ x_{k*S-1}$$

The Cascaded Information Bottleneck method relies on the principle that the relevant information can be progressively extracted from the time series with a cascade of successive bottlenecks sharing the same cardinality of bottleneck states but trained independently. The agglomerative information bottleneck algorithm is applied to each bottleneck successively, the first one being trained in the standard way while the next ones depend on the previous bottleneck states, as the following graph shows:

[2]One HMM per class to distinguish.

[3]In this context, the relevant information is defined as the information that the signal $x \in \mathcal{X}$ provides about another signal $y \in \mathcal{Y}$.

[4]for more details on the notion of mutual information, please refer to [15], [16]

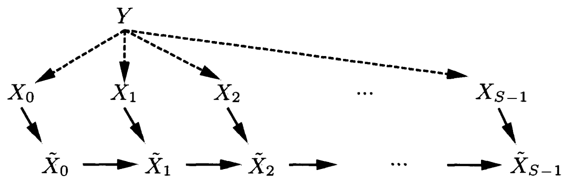

B. Extrapolation

The Cascaded Information Bottleneck method progressively extracts the relevant information from an input sample $X = [X_0, ..., X_{s-1}]$ by a recall on the successive components (X_0 for the first step of the cascade, $(\tilde{X}_{s-1}, X_s)$ for the other steps s). Each bottleneck is characterized by a hard mapping between: i) X_0 and $\tilde{X}_0$ for the first step, and ii) $(\tilde{X}_{s-1}, X_s)$ and $\tilde{X}_s$ for the other steps of the cascade. At each step s of the cascade, the algorithm looks for the equivalent $\tilde{x}_s$ given the input $(\tilde{x}_{s-1}, x_s)$ according to the hard mapping at step s (the equivalent $\tilde{x}_s$ satisfies the equation $p(\tilde{x}_s|(\tilde{x}_{s-1}, x_s)) = 1$).
It can happen that at a specific step s of the cascade, the pair $(\tilde{x}_{s-1}, x_s)$ for which we need to find the equivalent $\tilde{X}_s$ has never been encountered during the training process of this bottleneck. This pair is called an unseen pair. In the case of an unseen pair $(\tilde{x}_{s-1}, x_s)$ at step s, the cascade can a priori make no inference on $\tilde{X}_s$ because there is no preexisting default continuation of the cascade, due to the fact that the bottlenecks have been trained independently. In other words, for each pair $(\tilde{x}_{s-1}, x_s)$ which was not part of the training set data, $p(\tilde{x}_s|(\tilde{x}_{s-1}, x_s))$ is a priori undefined, whatever $\tilde{x}_s$ we take. For such cases, it is necessary to introduce a 'default' way leading from $\tilde{X}_{s-1}$ to $\tilde{X}_s$, i.e. we have to introduce an artificial identification of the bottleneck states which consists in matching out two bottleneck states (one at step $s - 1$ and one at step s). Therefore we apply a reorganisation of the bottleneck states at each possible step s (i.e. a one-to-one mapping of the bottleneck states at step $s - 1$ and the ones at step s which we call a permutation). For this purpose, we introduce the following measure $d_{(s-1,s)}$ allowing to directly compare the reorganised bottleneck states from step s with those from step $s - 1$. Let $\tilde{\mathcal{X}}_{s-1}$ (respectively $\tilde{\mathcal{X}}_s$) be the set of bottleneck states $\tilde{x}_{s-1}$ (respectively $\tilde{x}_s$) and $p(\tilde{x}_{s-1})$ (respectively $p(\tilde{x}_s)$) the empirical probability; for each permutation r of the bottleneck states $\tilde{X}_s$:

$$d_{(s-1,s)}(r) = -\sum_{\tilde{x}_{s-1} \in \tilde{\mathcal{X}}_{s-1}} p(\tilde{x}_{s-1}) \log \tilde{p}(\tilde{X}_s = r(\tilde{x}_{s-1}) | \tilde{X}_{s-1} = \tilde{x}_{s-1}) \quad (2)$$

Note that if the conditional probability $\tilde{p}(\tilde{X}_s = r(\tilde{x}_{s-1})|\tilde{X}_{s-1} = \tilde{x}_{s-1}) = 0$ then, by convention, $d_{(s-1,s)}(r)$ is ∞. The logarithm measures the unpredictability of the next case (i.e. the unpredictability of $\tilde{X}_s$ given $\tilde{x}_{s-1}$). We want to choose r to minimize that unpredictability and weight for the probability that the state $\tilde{x}_{s-1}$ actually happens (because there is no sense in penalizing a deviation if the state does not happen.). We call this permutation $R(s - 1, s)$.
The permutation of the bottleneck states that extracts the most similarity between bottleneck states at step $s - 1$ and those at step s is given by:

$$R(s-1, s) = \arg\min{}_r d_{(s-1,s)}(r) \quad (3)$$

We consider $R(s - 1, s)$ as the 'default' path between $\tilde{X}_{s-1}$ and $\tilde{X}_s$, i.e. as the criteria for extrapolating an unseen event at step s.

V. Application to the Recognition of Human-Robot Interaction Styles: Experiments

In this section we present an application of the Cascaded Information Bottleneck Method with real data: the automatic recognition of tactile interaction styles in the context of human-robot interaction. We conducted two series of trials, the first one under laboratory conditions and the second one in a school where several children could interact (one child at a time) freely with the robot. In all experiments the robot is the Sony Aibo and we focus on characterizing the tactile interactions according to two criteria, namely the *gentleness* and *the frequency of the interaction*. An interaction is classified as 'gentle' (respectively 'strong') if the participant strokes the robot gently, without signs of force (respectively with signs of force). The frequency of interaction is categorized into four classes S_i, $i = 0...3$, defined by their typical periodicity of interaction[5] T (in seconds): i) S_0: 'very low' ($T > 15$ seconds), ii) S_1: 'middle inferior' ($5 < T \leq 15$), iii) S_2: 'middle superior' ($1 < T \leq 5$), and iv) S_3: 'very high' ($T \leq 1$ second).

A. Implementation

1) Preprocessing: Each criterion (*gentleness* and *frequency of the interaction*) is studied independently. In each case, the time series studied is the quantitatively binned sum of the normalized sensors values[6] involved in the type of interaction.

2) Extra-conditions for the training: a) for the criterion 'gentleness', the algorithm does not learn null samples (i.e. samples made of null events only), b) for the frequency of interaction, the system deals only with samples whose first component is not null. In both cases, a sliding window proceeds on the sensor data time series.

3) Postprocessing: The postprocessing relies on a 'winner takes all' principle: The selected (winner state) is defined by $\arg\max_{y \in Y} p(y|\tilde{x}_{S-1})$.

B. Features of the trained cascade

The mutual information is 0.8 bit for the criteria gentle/strong and 1.9 bits for the frequency of the interaction. The conditional entropy $H(\tilde{X}_{s+1}|\tilde{X}_s)$ (Fig. 1) is globally decreasing over the cascade, pretty quickly, which suggests that a structure is progressively and rapidly emerging over the cascade. For the frequency of interaction, $H(\tilde{X}_{s+1}|\tilde{X}_s)$ has some small local peaks though, both at the very beginning

[5]The typical periodicity represents the elapsed time between two successive strokes of the robot.

[6]The robot's sensor data are updated every 32ms.

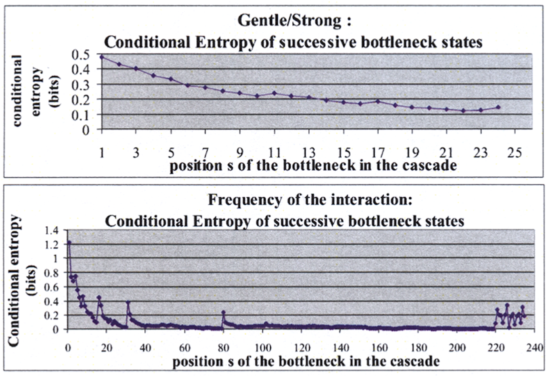

Fig. 1. Conditional entropy $H(\tilde{X}_{s+1}|\tilde{X}_s)$. There are four main parameters for the cascade: l (length of the input vector), k (length of the individual subsequences), S (length of the cascade), m (number of bottleneck states). For the frequency of interaction, $l = 472$ (equivalent to 15.1 seconds) , $k = 2$, $S = 236$, and $m = 6$. For the criterion gentle/strong, the corresponding parameters are: $l = 50$ (1.6 seconds), $k = 2$, $S = 25$ and $m = 4$.

of the cascade and at the very end[7], which suggest that at these steps s, the input data X_s may influence a bit more in the choice of next equivalent state $\tilde{X}_{s+1}$. This measure is correlated with the reorganisation measure for extrapolating $d_{s-1,s}(R(s-1,s))$ (equation (2) and equation (3)) which presents, respectively to each criterion of interaction, profiles similar to the conditional entropy with peaks positioned at the same place in the cascade (the mean of $d_{s-1,s}(R(s-1,s))$ is equal to, respectively, for Gentle/Strong, 0.037 bits, and, for the frequency of interaction 0.129 bits). In this study, the algorithm will extrapolate between step 5 and 24 (respectively 5 and 216) of the cascade for the gentleness (respectively frequency of interaction).

C. Experiments

The experiments aim at assessing statistically:

a) the soundness of the recognition of interaction styles by our algorithm, i.e.:
   i) for the criterion 'gentleness', whether a behaviour that has been classified as gentle (respectively strong) by a human is indeed going to be classified as gentle (respectively strong) by our algorithm,
   ii) for the frequency of interaction, whether a frequency of interaction that has been tagged by a human is indeed going to be correctly recognised by the algorithm.
b) the delay for the recognition of local events (i.e. short-term time scale events).

Importantly, the criterion 'gentle/strong' characterizes local events, and the algorithm should be able to recognise each specific event 'gentle' or 'strong' within a short delay. In contrast, the criterion 'frequency of the interaction' requires the algorithm to classify mid-term time scale events. This study deliberately focuses on such different criteria in order to show the flexibility of the algorithm.

1) Experimental setup under laboratory conditions: These trials are used as a first step in the statistical assessment of the soundness of the recognition of the interaction styles. They involve one participant at a time who is asked to interact with the robot for a few minutes in a *predefined way* which is one of the following:

- for the 'frequency of the interaction': only 'pure styles of interaction', i.e. one class[8] exclusively.
- for the criterion 'Gentle/Strong': In a first step, it is pure styles exclusively[9]. In a second step, the participant is asked to alternate gentle and strong behaviour and, just before generating the first event of the new class, he/she must name the style (i.e. "gentle" or "strong"). All the sessions are video recorded and this tagging enables to determine very precisely the transitions for a further measure of the delay of the recognition process.

2) Experimental setup in school: A further step in the validation of the algorithm is the testing with data obtained under natural situations of Human-Robot interaction. These experiments took place in a small classroom dedicated to the study, one child at a time being present in the room. Each child was invited to play freely for several minutes with the robot (the duration of play depended on the child's needs and abilities) in an unconstrained environment.

D. Measures

The experiments were all video-recorded and sensor data were stored. Note that the validation of the algorithm must be assessed offline but the recognition algorithm is designed to operate real time.

1) Samples excluding transitions from one class to another: The profile of the classification by the algorithm can be analysed with a confusion matrix which displays the probability distribution that events from class S_i are recognised by the algorithm as events of class S_i' ($i = 0$ or 1 for gentle/strong, $i = 0...3$ for the frequency of interaction).

2) Samples with transitions for the criterion gentle/strong: These samples enable us to test the ability of the algorithm to recognise a transition and reach, after a short transition phase, a new equilibrium phase. One can model this process by a temporal curve that would indicate the state of the system for a transition happening at time t_0. Three typical domains can be identified: for $t < t_0$ the curve is constant, indicating a stable state; from $t = t_0$, the curve's value alternates to indicate an hesitation between the two possible states (thus identifying a change in the behaviour observed); from $t = t_0 + \tau$ the curve would keep the same value (the new state). Ideally, the second phase should be very short (i.e. τ is very small). We will study three typical measures here: a) the number of transitions recognised by the algorithm; b) the time elapsed to

[7]Note that the small local peaks at the end of the cascade may reflect the importance of the last steps for distinguishing the classes S_0 and S_1.

[8]very low, middle inferior, middle superior, or very high.

[9]gentle or strong only.

reach the new equilibrium state, c) the ratio of errors made within this new equilibrium state. Note that a transition will be considered broadly as either a transition from a gentle (respectively strong) behaviour to a strong (respectively gentle) one, or from a state where no classification occurred (i.e. no interaction occurred during the past 1.6 seconds) to gentle or strong.

3) Samples with hybrid behaviours for the frequency of interaction: Because this criterion is based on a mid-term time scale analysis, some samples generated in school can be hybrid, i.e. contain a mix of features from different classes. In order to encapsulate hybrid behaviours, the human classifies the behaviours on a 'two choices' basis, i.e. he/she can select the two styles characterising the hybridity. In this case, the algorithm's classification is successful if it agrees with one of the two choices made by visual inspection.

Practically, the video and graphs of the temporal global variable are first manually tagged. In a second step, the classifications S_i resulting from the manual tagging are compared with the classifications S_i' made by the algorithm.

VI. Application to the Recognition of Human-Robot Interaction Styles: Results

We present the results for each criterion (gentleness and frequency of the interaction) successively. Note that here we will refer to the samples of data that were classified without using the extrapolation, i.e. the samples that contained no unseen cases at any step of the cascade, as *samples classified without extrapolation*. In contrast, the samples of data that required an extrapolation at one or more steps of the cascade, i.e. the samples for which there were unseen cases to extrapolate (i.e. cases that had not been encountered during the training phase of the algorithm), will be referred to as *samples classified with extrapolation*.

A. Criterion: Gentle/Strong

1) Training set of data: The $20,018$ samples used for the training were classified by the algorithm with an overall success of 97.82% and, respectively, for gentle and strong, 96.83% and 98.81%.

2) Samples excluding transitions (cross-validation): They constitute 1 hour 2 minutes 49 seconds of interaction. $100,111$ samples have been classified with a ratio of success for correct classification of 0.948. 97.7% of samples were classified without extrapolation with 95.22% of success while the samples classified with extrapolation (3.3%) were well classified in 75.54% of cases which, considering that it results from an extrapolation, is quite a good result. Note that the parameters of the Cascaded Information Bottleneck Method were chosen in such a way to have a good balance between the extrapolation and the precision, which is reflected here in the low percentage of cases extrapolated.

3) Samples with transitions under laboratory conditions (cross-validation): The four runs constitute 19 minutes and 40 seconds of interaction to analyse. They contain $53,192$ samples to classify and 0.01% of the samples were not classified because they could not be extrapolated by the algorithm[10]. 212 transitions were to be recognised, 99.1% of which were indeed well classified by the algorithm[11] with an average delay of 0.17 seconds. The cumulative probability distribution of the delay is displayed in Fig. 2. The curve grows very rapidly, thus showing that most of the delays are very small. Transitions recognised without any delay occur particularly in the case of a transition from no event to classify to any event to classify. The longest delay is 2.05 seconds, which we consider very acceptable for human-robot interaction kinesics. The average error ratio in the equilibrium phase is 0.02 and the cumulative probability distribution is displayed in Fig. 3. Here again, the curve grows rapidly and shows that the probability of the highest error ratio is very low and remains acceptable for real human-robot interaction.

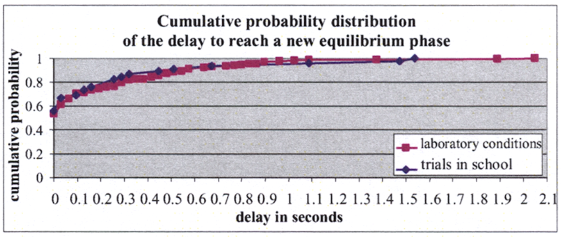

Fig. 2. Cumulative probability distribution of the delay for recognising the transition. We display the probability that an event is recognised within (less or equal) n seconds for a given n. The delay corresponds to the length of the transition phase when a transition occurs.

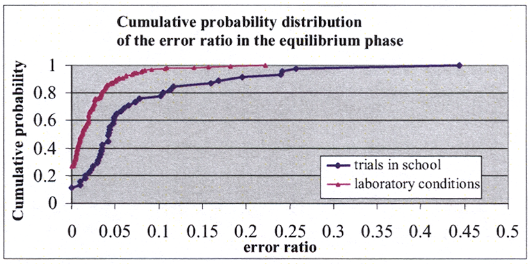

Fig. 3. Cumulative probability distribution of the error ratio for the equilibrium phase. The ratio measures the number of errors of classification made during a phase of equilibrium divided by the number of samples to classify during this phase. The figures displayed give, for a given r, the probability that the error ratio is inferior or equal to r.

4) Samples generated by the children in the school (cross-validation): Videos from five different children were analysed, which constitute 12 minutes and 52 seconds of interaction. These runs contain $6,660$ samples to classify: 97.49% of these samples have been classified by the algorithm. These samples contain 45 transitions. 91.1% of these transitions were indeed well classified by the algorithm within an average delay of

[10] these samples had to be extrapolated outside the range of steps considered for the extrapolation.

[11] A transition is considered as wrongly classified if the transition phase is very long compared to the new equilibrium phase.

0.17 seconds. The cumulative probability distribution of the delay is represented in Fig. 2. The curve grows very rapidly, thus showing that most of the delays are very low. Transitions recognised without any delay occur, and, at the far end, the highest delay is 1.54 seconds, which is very acceptable for human-robot interaction kinesics. The mean error ratio in the equilibrium phase is 0.1 and the cumulative probability distribution of this ratio is displayed in Fig. 3. Here again, the curve grows rapidly. It is worthy of note that the highest value obtained is 0.44 and the second one is much lower (0.26) which indicates that the first highest value can be seen as an extraordinary case. Looking at the sequential classification of the results, it appears that this highest error ratio was obtained while a child interacted in a very instable way that is, within 1.76 seconds three successive transitions were observed that are 1) no event to gentle (gentle phase lasted 1.37 seconds), 2) gentle to strong (the phase with strong style lasted only 0.26 seconds), 3) strong to gentle. It is the strong phase, after the transition from gentle to strong behaviour that was recognised with the highest error ratio (0.44), but it lasted for such a short time that it is not really a concern here (0.26 seconds is very low compared to the typical time for human-robot interaction which usually lasts a few seconds). Therefore, we can consider to omit this highest value in the probability distribution and looking at the resulting values, the results are good and comparable to the results obtained in the laboratory.

B. Criterion: Frequency of the interaction

1) Training set of data: It constitutes 36 minutes 34 seconds of interaction and contains 4,865 samples to classify (respectively, 450 for S_0, 1,208 for S_1, 1,484 for S_2 and 1,723 for S_3). 99.98% of these samples are well classified; the ratio of success specific to each class is displayed in Fig. 4.

	S'0	S'1	S'2	S'3
S0	**1**	0	0	0
S1	0.0008	**0.9992**	0	0
S2	0	0	**1**	0
S3	0	0	0	**1**

Fig. 4. Confusion Matrix for the training set. The ratio is the one among events from type S_i. S_i represents the real class and S'_i the recognised class, $0 \leq i < 4$.

2) Samples generated under laboratory conditions (cross-validation): They constitute 51 minutes 44 seconds of interaction and contain 5,395 samples to classify (respectively 1,017 for S_0, 855 for S_1, 1,933 for S_2 and 1,590 for S_3) 91.16% of which were classified with an overall ratio of success of 0.922. 99.4% of the samples not extrapolated were well classified, and 76.41% of samples classified through extrapolation were well classified. Fig. 5 displays the confusion matrices.

3) Samples generated by the children in the school (cross-validation): Three runs of interaction were used for the validation of the frequency of interaction in a real situation, from three different children. They constitute 14 minutes 41 seconds of interaction and contain 5,288 samples to classify. 91% were classified (including 26.81% that had to be extrapolated) and

No Extrapolation	S'0	S'1	S'2	S'3
S0	**1**	0	0	0
S1	0	**0.972**	0.028	0
S2	0	0	**0.999**	0.001
S3	0	0	0.006	**0.994**

Extrapolation	S'0	S'1	S'2	S'3
S0	**1**	0	0	0
S1	0.115	**0.864**	0.022	0
S2	0.083	0.146	**0.768**	0.003
S3	0	0	0.368	**0.632**

Fig. 5. Confusion Matrices for pure sets of data for, respectively, non extrapolated and extrapolated data. Non extrapolated samples are samples which were classified without the need to use the extrapolation, because none of the cases were unseen cases (relatively to the training set samples). The results for those samples are provided in the table with mention No extrapolation. On the contrary, extrapolated samples are samples that used the extrapolation at least once in the cascade (those samples contained at least one unseen case in the cascade, i.e. a case that had not been encountered during the training). The results for those samples are provided in the table with the mention Extrapolation. See Fig. 4 for more details on the notion of confusion matrix.

93% were classified correctly. Among samples classified with no extrapolation, the ratio of success for a sound classification was 0.96. while for samples classified with extrapolation, it was 0.84.

VII. Discussion and Future Work

The algorithm has proven sound for the recognition of the two criteria of interaction. Concerning the criterion gentle/strong, results show that the two classes are well recognised and the delays very acceptable for human-robot interaction. The extrapolation works well, which shows the capability of the system to make a sound decision in case of unseen events. These results can be compared with a previous study of ours where we used Self-Organizing Maps to classify this criterion of interaction [4], whereby the average delay to recognize transitions was much higher and the postprocessing required more effort.

Importantly, one might wish to define the styles slightly differently to the definition given here, such as, for instance, focusing on more details (in order to describe substyles for instance). This can be easily done by adjusting relevant parameters, mainly the number of bottleneck states, the binning and the training sets which condition the learning.

The algorithm has also proved very capable of classifying real data over a mid-term time scale (cf. the criterion frequency of the interaction) which illustrates the ability of the method to make a powerful exploitation of an existing temporal structure not only of short-term time scales but also mid-term ones. This ability is empowered by the use of different bottlenecks (thus different mappings) over the cascade. In contrast, as explained in the section on Related Work, with HMMs the mapping would be the same all over the time series, and, by trying to squeeze all temporal information into one flat transition structure, it might actually prevent HMMs from an efficient making use of an existing temporal structure of the data. This hypothesis should be investigated in future work which will include a comparison of our method with HMMs in these scenarios. The problem with a cascade of bottlenecks trained independently could be here that the system has too many degrees of freedom and could overlearn. The extrapolation with the measure that we have introduced is a first step in the control of the degrees of freedom of the system. In addition, the overlearning can be tightly controlled by penalizing the intake of novel information. For this, we would have to move from the agglomerative model where $\beta = \infty$ to a model with a finite β that would control the information intake per step. This shows how the Cascaded Information Bottleneck method is transparent and gives fine-grained control over how much and what new information is taken at which step in the cascade.

This method is designed for real-time use during natural human-robot interaction and little research had been done so far on real-time recognition of tactile interaction styles. Salter et al. 's adaptation algorithm [11] was a first important step towards real adaptation. Yet, this system did not learn its own categorisation, which was completely described by a hand-tuned decision tree. In the present study, the recognition and the decision are made algorithmically, after a real learning phase and a capacity to extrapolate unseen events, with very small delays. Furthermore, our method is very easy to use and can be tuned easily to adapt to other criteria of interaction.

VIII. Conclusion

In this paper, we have presented a novel method for time series analysis for detecting interaction styles in the context of Human-Robot Interaction. This method, namely the Cascaded Information Bottleneck Method relies on a cascade of bottlenecks trained independently, the first one being trained in a standard way [3] while the next ones depend on the previous bottleneck states. This notably facilitates a powerful exploitation of the temporal structure of the data. Besides, a structure progressively emerges through the cascade and we introduced a measure to extrapolate unseen cases.

We have applied our method to real-time recognition of human-robot interaction styles, in a detailed case study, by implementing the algorithm for real interactions with a real robot. The testing of the method had to be done offline, i.e. after the interactions had taken place, but the algorithm is designed to operate real time in order to enable real-time adaptation of robots to the interaction styles.

We have shown the soundness of the method through extensive experiments, using successively samples of data generated under laboratory conditions and samples from natural situations of child-robot interaction in a school for children with autism. The algorithm was able to recognize short term events very well within and average delay of 0.17 seconds (the highest delay being 2.07 seconds). It was also able to recognise mid-term time scale events very well (the percentage of events correctly classified was 92% under laboratory conditions and 93% with data from the child-robot interactions).

This study has shown the soundness of the method for pattern recognition and illustrated its capability of time-filtering on real data. Besides, the method is transparent and enables a fine-grained control over how much and what new information is taken at which step of the cascade. Finally, this method is entirely generic for applications with socially interactive (humanoid and non-humanoid) robots.

Our own future work will focus on the application of the method in autism therapy where we find a strong need for socially adaptive robots. The ability of a robot to classify in real time human-robot interaction styles is a first step towards the challenging goal of enabling an autonomous robot to influence positively children's interaction styles to guide him/her progressively towards different therapeutically relevant levels of interaction.

Acknowledgments

Dorothée François is supported by a research scholarship of the University of Hertfordshire. The work described in this paper was partially conducted within the EU Integrated Project RobotCub (Robotic Open-architecture Technology for Cognition, Understanding and Behaviours) and was partially funded by the European Commission through the E5 Unit (Cognition) of FP6-IST under Contract FP6-004370. The authors would like to thank the children who participated in these trials, their parents and their teachers.

References

[1] "Aurora project," 2008, http://www.aurora-project.com/.
[2] C. E. Shannon, *The mathematical theory of communication*. The University of Illinois Press, Urbana, 1949.
[3] N. Tishby, F. C. Pereira, and W. Bialek, "The information bottleneck method," in *Proc. of the 37-th annual Allerton Conference on Communication, Control and Computing*, 1999, pp. 368–377.
[4] D. François, D. Polani, and K. Dautenhahn, "On-line behaviour classification and adaptation to human-robot interaction styles," in *Proc. 2nd ACM/IEEE International Conference on Human-robot Interaction (HRI07)*, 2007, pp. 295–302.
[5] T. Shibata, K. Wada, T. Saito, and K. Tanie, "Human interactive robot for psychological enrichment and therapy," in *Proc. AISB'05 Symposium on Robot Companion Hard Problem and Open Challenges in Human-Robot Interaction*, 2005, pp. 98–109.
[6] P. Marti, F. Fano, V. Palma, A. Pollini, A. Rullo, and T. Shibata, "My gym robot," in *Proc. AISB'05 Symposium on Robot Companion Hard Problem and Open Challenges in Human-Robot Interaction*, 2005, pp. 64–73.
[7] W. D. Stiehl, J. Lieberman, C. Breazeal, L. Basel, L. Lalla, and M. Wolf, "The design of the huggable: A therapeutic robotic companion for relational, affective touch," in *AAAI Fall Symposium on Caring Machines: AI in Eldercare*, 2006.
[8] B. Scassellati, "Quantitative metrics of social response for autism diagnosis," in *Proc. 14th IEEE Int. Workshop on Robot and Human Interactive Communication (RO-MAN)*, 2005, pp. 585–590.
[9] T. Kanda, H. Ishiguro, T. Ono, M. Imai, and R. Nakatsu, "Development and evaluation of an interactive humanoid robot "'robovie"'," in *Proc. IEEE Int. Conf. on Robotics and Automation (ICRA2002)*, 2002, pp. 1848–1855.
[10] T. Salter, K. Dautenhahn, and R. te Boekhorst, "Learning about natural human-robot interaction," *Robotics and Autonomous Systems*, vol. 54(2), pp. 127–134, 2006.
[11] T. Salter, F. Michaud, D. Lee, and I. P. Werry, "Using proprioceptive sensors for categorizing human-robot interactions," in *Proc. of the 2nd Int. Conference on Human-Robot Interaction HRI07*, 2007, pp. 105–112.
[12] A. Derakhshan, F. Hammer, Y. Demazeau, and H. H. Lund, "Mapping children and playgrounds into multi-agent systems," in *Proc. 11th Int. Symposium on Artificial Life and Robotics (ISAROB)*, 2006.
[13] D. Kim, J. Song, and D. Kim, "Simultaneous gesture segmentation and recognition based on forward spotting accumulative hmms," *Pattern Recognition*, vol. 40(11), pp. 3012–3026, 2007.
[14] C. Lee and Y. Xu, "Online, interactive learning of gestures for human/robot interfaces," in *Proceedings of the IEEE International Conference on Robotics and Automation*, 1996, pp. 2982–2987.
[15] J. P. Crutchfield, *Information and its Metric*. Nonlinear structures in physical systems- Pattern Formation, Chaos and Waves, L.Lam and H. C. Morris Eds. Springer Verlag, pp 119-130, 1990.
[16] T. M. Cover and J. A. Thomas, *Elements of Information Theory*. Wiley-Interscience, New York, 1991.
[17] N. Slonim and N. Tishby, "Agglomerative information bottleneck," in *Proc. of Neural Information Processing Systems (NIPS 99)*, 1999, pp. 617–623.

SECOND PROOFS OF ARTICLE IN PRESS

A long-term study of children with autism playing with a robotic pet

Taking inspirations from non-directive play therapy to encourage children's proactivity and initiative-taking

Dorothée François[1], Stuart Powell[2] & Kerstin Dautenhahn[1]

[1]Adaptive Systems Research Group, School of Computer Science, University of Hertfordshire/[2]School of Education, University of Hertfordshire

This paper presents a novel methodological approach of how to design, conduct and analyse robot-assisted play. This approach is inspired by non-directive play therapy. The experimenter participates in the experiments, but the child remains the main leader for play. Besides, beyond inspiration from non-directive play therapy, this approach enables the experimenter to regulate the interaction under specific conditions in order to guide the child or ask her questions about reasoning or affect related to the robot. This approach has been tested in a long-term study with six children with autism in a school setting. An autonomous robot with zoomorphic, dog-like appearance was used in the studies. The children's progress was analyzed according to three dimensions, namely, Play, Reasoning and Affect. Results from the case-study evaluations have shown the capability of the method to meet each child's needs and abilities. Children who mainly played solitarily progressively experienced basic imitation games with the experimenter. Children who proactively played socially progressively experienced higher levels of play and constructed more reasoning related to the robot. They also expressed some interest in the robot, including, on occasion, affect.

Keywords: Human–Robot Interaction, Robot-Mediated Therapy, Robot-Assisted Play, Non-Directive Play Therapy, Assistive Technology, Autism, Children

1. Introduction

This study is part of the Aurora Project, (Aurora, 2009) an ongoing long-term project investigating the potential use of robots to help children with autism overcome some of their impairments in social interactions (Dautenhahn & Werry, 2004, 2000).

Interaction Studies 10:3 (2009), **324–373**. DOI 10.1075/is.10.3.04fra
ISSN 1572–0373 / E-ISSN 1572–0381

Autistic spectrum disorders can appear in various degrees and refer to different needs and abilities (Powell, 2000; American Psychiatric Association, 1994). Detailed diagnostic criteria for autistic spectrum disorders are provided in the Diagnostic and Statistical Manual of Mental Disorders (1994). The main impairments highlighted by the National Autistic Society[1] are: impairments in communication, social interaction and imagination. As a consequence, children with autism often seem to operate in a world of repetitive patterns and some of them tend to restrict play to solitary play. Besides, it can be argued that children with autism have a relative potential for play but often encounter obstacles to actualize this potential, the causes of which are still under investigation. Difficulties in socio-emotional inter-subjectivity, joint attention and theory of mind (compare e.g. Baron-Cohen et al. (1985); Hobson (1993); Baron-Cohen (1997)) impair interactions in general and, more specifically, imply a lack of spontaneous and social reciprocity during play. Those impairments, in addition to the potential deficits in higher order representation, may explain the difficulties encountered in symbolic and pretend play (Chaillé & Silvern, 1996).[2] Yet, play is a vehicle for learning (Chaillé & Silvern, 1996). Through play, children can develop skills in many areas (e.g. logical memory and abstract thought, communication and social skills). Moreover, play is a medium for self expression (Boucher, 1999). From the perspective of this study that aims at supporting robot-assisted therapy for children with autism, we thus decided on an emphasis on play whereby the robot should facilitate play and adapt to each child's needs and abilities.

The use of robots for robot-assisted play and therapy is a growing area of research (see section on 'Related Work'). Until now, many approaches of robot-mediated play and therapy for children with autism have mainly explored the use of specific games, such as imitation (Robins, Dautenhahn, Boekhorst, & Billard, 2005) or chasing games (Werry & Dautenhahn, 1999) and only recently started to involve the experimenter in the play sessions, qualifying his/her role as "passive participant" (Robins, Dautenhahn, 2006). The study presented in this paper shows a different perspective on robot-mediated therapy, which is not primarily task-oriented. It draws inspiration from non-directive play therapy (Axline, 1946, 1947; Ryan, 1999; Josefi & Ryan, 2004) and, importantly, expands and *formalizes* the role of the experimenter in robot-assisted play. In this novel approach, the experimenter strongly encourages the child's proactivity and initiative-taking with respect to the choice of play, the rhythm of play and verbal communication. While a task-oriented approach might expect the child to complete a specific task, such as, for instance, performing imitative movements, our approach enables the child to proactively experiment with various situations of play, from simple exploration of the robot's features and capabilities to more complex situations of play. Those situations can, for instance, involve an understanding of the notion of causality or the ability to take on a specific role in play. Furthermore, at any moment, the child can appeal to the experimenter's participation, thus enabling the child to experience triadic play.

Besides, beyond inspiration from non-directive play therapy, the approach presented in this paper introduces a *regulation process*. This process notably enables the experimenter to regulate the interaction in order to guide the child towards other play styles when needed or modify slightly the rhythm of play if she feels the child is "standing still". The study presented in this paper explores the potential of this novel methodological approach for robot-assisted play through a case-study evaluation of a long-term study with six children with autism. This study should be regarded as a preliminary exploration of the feasibility of such a technique in the context of robot-mediated therapy for children with autism. Several research questions are addressed:

a. Does such an approach of robot-mediated therapy, inspired by non-directive play therapy, help the child experience higher levels of play and enable him/her to develop new play skills?
b. Does this approach encourage the child to play socially?
c. Might this approach be appropriate for children who play solitarily and speak mostly by using onomatopoeia?[3] Might it help him/her experience social play? If not, what might be the additional requirements necessary for such experience?

2. Non-directive play therapy

This section summarizes the core ideas of non-directive play therapy as mainly developed in Axline (1947) and explained and illustrated by case studies in Ryan & Wilson (1996).

Non-directive play therapy has its roots in Rogerian client-centred therapy with adults (Rogers, 1976), adapted to child therapy with a focus on play as the principal medium of communication (in contrast to verbal exchange). Rogerian theory[4] relies on the idea that all human beings have a drive for self-realisation; it means that any human being tends to develop towards maturity, independence and self-direction. The individual needs to completely accept himself/herself as well as be accepted by others.

In non-directive play therapy, the child, rather than the therapist, chooses the type of play and the activity in general in the playroom. This contrasts with other play interventions. We shall cite Axline who primarily developed the method of non-directive play therapy (Axline, 1947): "Non-directive play therapy is not meant to be a means of substituting one type of behaviour, that is considered more desirable by adult standards, for another 'less desirable'. It is not an attempt to impose upon the child the voice of authority that says 'You have a problem. I want you to correct it.'" Few limitations in the behaviour of the child are set, which refer to safety and security reasons.

A relationship is progressively built up between the child and the therapist. This relationship enables the child to share his/her inner world with the therapist and, "by sharing, (the child) extends the horizons of both their worlds" (Axline, 1947). Ryan et al. state that this relationship, with the help of the therapist, progressively facilitates the child to choose freely the feelings he/she wishes to focus on as well as the way in which he/she wants to explore them (Ryan & Wilson, 1996). Three mediums may be used for communicating these feelings: action, language and play.

The therapist participates in the therapy. She observes, listens to and answers the child. The therapist is reflecting the child's feelings or emotionalized behaviours in order to help him/her build a better understanding of himself/herself. The therapist's role has been characterized by eight basic principles set out by Axline (Axline, 1947), see Fig. 1.

Note that in the study presented in this paper the experimenter was not trying to engage in therapy; the study only drew inspiration from non-directive play therapy, thus the context may be a therapeutic one, but the experimenter, a human–robot interaction researcher, was not behaving exactly like a therapist. The experimenter was not applying strictly the eight principles set out by Axline (Axline, 1947), see Fig. 1. She very much drew inspiration from Axline's principles 1, 2, 3, 5 and 8, but she was not dealing with the fourth one; and, concerning principles 6 and 7, she was considering these principles with more flexibility. It is worthy of note here that this study is a first step towards a proof-of-concept and required significant robotics expertise; however, in future, play therapists may use this approach.

3. Related work

3.1 Non-directive play therapy for children with autism

Non-directive play therapy has been largely used for children and adolescents with a wide variety of emotional and behavioural problems (Ryan, 1999; Ryan & Needham, 2001). Only recently have researchers started to investigate the feasibility of such techniques with children with autism. A pioneering case study was presented in 2004 in Josefi & Ryan (2004) involving a 6-year-old-boy with severe autism. The child attended 16 non-directive play therapy sessions of an hour a 5-month period in the child's special school. The room was empty except for specific materials selected for their "expressive, imaginative, relaxing and interactive properties". Results were analysed both qualitatively and quantitatively. Results showed an increase in the child's autonomy and initiative-taking and the child developed an attachment to the therapist. According to Josefi et al. (Josefi & Ryan, 2004), it was shown that non-directive play therapy itself may provide children with autism with the basis[5] for therapeutic progress as stated in play literature (Axline, 1947). Also, the child's concentration increased and his repertoire of play

1. "The therapist must develop a warm, friendly relationship with the child, in which good rapport is established as soon as possible."
2. "The therapist accepts the child exactly as he is."
3. "The therapist establishes a feeling of permissiveness in the relationship so that the child feels free to express his feelings completely."
4. "The therapist is alert to recognize the *feelings* the child is expressing and reflects those feelings back to him in such a manner that he gains insight into his behavior."
5. "The therapist maintains a deep respect for the child's ability to solve his own problems if given an opportunity to do so. The responsibility to make choices and to institute change is the child's."
6. "The therapist does not attempt to direct the child's actions or conversation in any manner. The child leads the way; the therapist follows."
7. "The therapist does not attempt to hurry the therapy along. It is a gradual process and is recognized as such by the therapist."
8. "The therapist establishes only those limitations that are necessary to anchor the therapy to the world of reality and to make the child aware of his responsibility in the relationship."

Figure 1. Eight basic principles set out by Axline for practice of non-directive play therapy: quotations from Axline (1947)

expanded over the sessions. The games involved progressively more joint attention and direct social interaction and verbal communication with the therapist increased; symbolic play emerged more and more verbally with the therapist.[6] However, repetitive and obsessive behaviours were not considerably reduced. As a conclusion, Josefi et al. (2004) stated that non-directive play therapy with children with autism may be complementary to behaviour therapy, non-directive play therapy is likely to be more efficient in the child's gaining autonomy, taking initiative, showing joint attention and developing social and symbolic play, while behaviour therapy could be more efficient in reducing ritualistic and obsessive behaviours.

3.2 Robot-mediated therapy and education

Robot-mediated therapy is an area of research in assistive and rehabilitation robotics that aims at using robots in the therapy of patients in a variety of domains, e.g. in stroke rehabilitation (Loureiro et al., 2003). Robot-mediated therapy, and in particular the use of robot-assisted play in therapy or education, is a growing research field. It has been shown that robots, compared to simple toys, elicit a range of behaviours in children with autism that are more desirable in the light of encouraging and/or teaching children with autism social behaviour and communication. Werry and Dautenhahn (Werry, Dautenhahn, & Harwin, 2001; Dautenhahn & Werry, 2004) showed that children with autism exhibited more eye gaze and more attention directed towards an autonomously operating mobile robot compared to a non-robotic toy. Later, Stanton et al.'s studies compared interactions of children

with autism with an Aibo robot (Sony) and a simple mechanical toy (Stanton et al., 2008). Their results show that the children spoke more words to the robot and more often showed certain behaviours towards the Aibo including verbal engagement, reciprocal interaction, and authentic interaction. Such comparative studies provide the main motivation of our approach to use robots (and not other non-robotic toys) to investigate their potential in the therapy and education of children with autism.

A fully comprehensive review of the literature would go beyond the scope of this paper, and we therefore focus below on selected research that is particularly relevant to the present work.

Long-term studies with the seal robot Paro have shown that specific everyday life situations exist in which human–robot interaction can have a positive effect on the well-being of human beings (Shibata et al.,2005); they may even be a significant factor of performance in therapy[7] (Marti et al., 2005).

Outside the therapeutic context, in the broad field of child–robot interaction, Tanaka et al. led a long-term study in a school in order to identify principles for realizing long-term interaction (Tanaka et al., 2005, 2006). This study notably showed that the children's views of the robot evolved: they progressively considered the robot (in this case, the humanoid robot QRIO) as a peer rather than as a toy.

Within the Aurora Project, Robins et al. carried out long-term studies analyzing, on the one hand, the role of the robot as a mediator (Robins, Dautenhahn, & Dubowski, 2005) and, on the other hand, the role of the experimenter (Robins & Dautenhahn, 2006) which they described as that of a "passive participant" who responds to the children if they initiate interaction with him/her. There was no autonomous reaction from the robot to the child's interactions in their study. Moreover, child–robot interaction situations taking place during these trials were mainly concerned with encouraging imitation of gestures (position or movement of arms and legs). In Robins et al.'s experiments, children interacted with a remotely controlled robotic doll by imitation of gestures.[8]

In different studies, Werry et al. encouraged free-play with a mobile rectangular autonomous robotic platform, Labo-1, equipped with infrared sensors (Werry & Dautenhahn, 1999; Werry, Dautenhahn, Ogden, & Harwin, 2001). The play situations were approach and avoidance games whereby turn-taking emerged from the child–robot interactions (Dautenhahn, 2007). The experimenters did not take part in the games; they only responded to the child when the child initiated communication or interaction with them (Dautenhahn & Werry, 2002).

Outside the Aurora Project, Kozima et al. used a small dancing creature-like robot, Keepon, in a long-term study with children with autism, in relatively unconstrained conditions (Kozima et al., 2005). Keepon was manually controlled by the experimenter who was not part of the trials. Children and carers were involved in the trials which highlighted the role of Keepon as a pivot in triadic interaction by facilitating the emergence of joint attention. Another study conducted by Duquette et al. (2007) showed the potential of the robot Tito to elicit shared focused attention[9] (visual contact, physical proximity) in a large range of imitation games. This study also pointed out the impact of the robot in imitation games and showed its potential to foster imitation of facial expression but also, in this specific context, its limits for encouraging e.g. imitation of gestures.

These results reinforce the idea that child–robot interaction may be valuable for children with autism with respect to being a medium towards possible social interactions. It also shows the relevance of investigating new approaches in how to design and conduct robot-assisted play for children with autism.

4. Method

4.1 Participants

All the children taking part in the experiments have a diagnosis of autism and are from the same school based in the UK. This school welcomes children between 4 and 11 years old with moderate learning difficulties. In particular, an Autism Base provides extra care and a specific education program for children with autism to start within the school. When the child gets older or when he/she has made sufficient progress (especially if he/she has improved in social skills) he/she can be integrated in a more general class for children with specific needs and abilities including children with autism. Six children were selected by the teachers to take part in the current study. For purposes of clarity and simplicity, a consistent naming of the children will be used in the whole paper, starting with A and then, alphabetically, in order of appearance in the text.

Two boys from the Autism Base, Child A (seven years old) and Child B (eight years old) were invited to take part in the experiments. Both of them find it hard to express themselves verbally and their behaviour often includes using onomatopeia and repetitive gestures. According to the teachers, Child A often shows apprehension towards dogs and doors and Child B has a fascination for computers. Child C took part in the experiments who is a seven-year old girl. During the experiments, she was part of the Autism Base but in the process of being integrated into another class with children with moderate learning difficulties including children with autism. She therefore started to follow part-time the education program of this class and the rest of the time stayed in the Autism Base. She masters verbal communication pretty well and teachers describe her behaviour as proactively social, as far as play at playtime is concerned.

Three older children took also part in the experiments. All of them are integrated in classes for general moderate learning difficulties. Child D, ten years old, is described by his teacher as a solitary child. In the classroom the position of his desk, fairly isolated from the others, gives him his 'own' space. Child D understands pretty well when one addresses him verbally but mostly speaks by using onomatopeia. At school, he often uses the computer to do exercises, especially exercises on words and writing. Two other children, Child E, ten years old and Child F, nine years old, took also part in the study. They communicate verbally and are not described as solitary children.

Note, other details, such as mental age of the children, were not available. The study was carried out with approval of the University of Hertfordshire Ethics Committee. The parents of all the children who took part in this study gave written consent, including permission to videotape the children and utilize photos in publications.

4.2 Artifact

The main artifact is a white robotic mobile autonomous dog, the Sony Aibo ERS-7 (Fig. 2). It is equipped with a great variety of external sensors, and particularly, five tactile sensors: the head sensor, the chin sensor and the three back sensors. Aibo's control programming was achieved using URBI (Baillie, 2005). A laptop endowed

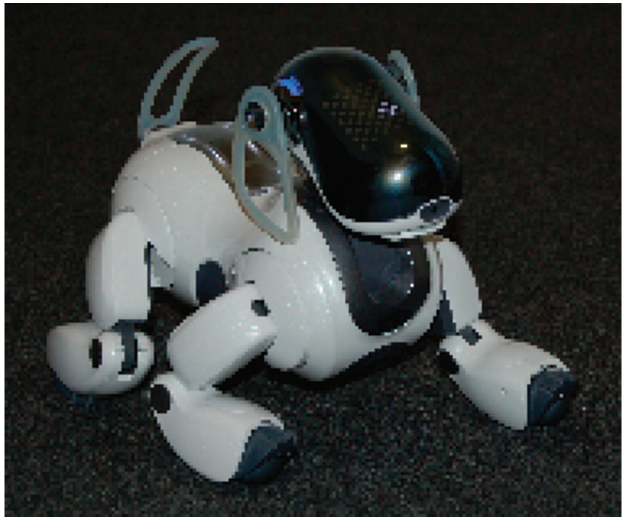

Figure 2. Aibo ERS-7. Aibo ERS-7 weights approximately 1.65kg and measures approximately 180(w) x 278(h) x 319(d) mm

the robot with specific behaviour-modes through a wireless connection. Once the robot had been endowed with a specific behaviour-mode, it reacted autonomously to the activation of its sensors.

4.3 Procedures and measures

4.3.1 *Procedures*

Experimental setup. The experiments took place once a week in the school. Each child participated in a maximum of ten sessions. Not every child could take part in 10 sessions because some of them may have been away for a day. Note that an exception was made for one child who showed some apprehension towards the robot: for him, experiments were stopped after 5 sessions and only restarted in the last session when he proactively came to the trial.

The rooms used for the experiments changed several times due to circumstances at the school (Fig. 3). In each case, the child could have encountered possible distractive objects, like toys or mirrors (Fig. 4). Thus, these experiments took place in a context of possible distraction.

Each trial involved one child with autism, the experimenter[10] and sometimes another researcher from the Aurora project with whom the children were familiar. The latter helped the experimenter film the trials and occasionally took part in a verbal communication process by answering a child's question directly addressed to her.

The duration of the sessions was variable. The child was free to play as long as he/she wanted with the following restrictions: (i) the upper limit of time was 40 minutes (so that the child did not miss too much of his/her courses at the school); (ii) if the child had an obligation due to his/her schedule, then the session was shortened.

The Aibo robot was programmed in order to show simple behaviours, tailored progressively by immersion according to each child's needs and abilities. Note that 'tailored by immersion' means here that the repertoire of appropriate robot behaviours with respect to each child's specific needs, abilities, dislikes and preferences was progressively refined as the experiments progressed. The mapping between the sensors and the reactions of the robot (also called behaviour-mode) could therefore vary from one session to the other and also during a session in order to meet as closely as possible the needs, abilities and demands of the child at a given moment (Fig. 5 includes examples of the robot's behaviours). The robot reacted autonomously to the activation of its sensors, with respect to the specific behaviour-mode it had been endowed with. The switch between various behaviour-modes was done manually by the experimenter through a wireless connection with a laptop. The laptop was located in the same room as the children, and thus constituted an additional source of distraction for the children.

Room	Description	Dimensions	Furniture in the room	Objects in the room
R 1	Small room	Approx. 10feet * 8feet	-small longitudinal window on the very top (children can't see through it), -cupboard, -low rectangular table, -2 children's chairs, -decoration on the wall (a clown's head drawn on a paper board).	Regular objects: - game with individual letters to form words, reflective blue metallic support, - coloured cubes (25mm*25mm) - rectangular paperboard 3D decoration, 1m*30cm*20cm , vertically in a corner. On occasion: man's like face drawn on a paperboard that children could hold in front of their face.
R 2	Small room in the Autism Base	Approx. 10feet * 12feet	-big window on a wall, -second internal window (semi-transparent, semi-reflective) with view on another classroom; -vertical mirror, children can see their whole body by reflection -shelves on the very top, children can't access -table & small chairs (session8 only)	- games in open boxes on the shelves (e.g. a doll); children can see them but can't access them.
R 3	Large meeting Room: library, kitchen and living room corners. Experiments took place in the living room corner.	-room: Approx. 35feet * 40feet; -living room corner, approx. 10feet * 12feet	-Large windows on two walls -2 sofas made of joint comfortable chairs -4 comfortable additional chairs -rectangular dinner table, 6 chairs -2 low coffee tables -shelves (at the entrance) -kitchen corner	-magazines on the coffee table -on the shelves, objects such as cloth samples in open boxes -small calculator -small alarm clock
R 4	Classroom; experiments took place in the library corner	-room: Approx. 30feet * 30feet; -library corner: approx. 10feet * 7feet	Library corner: -2 shelves separating the library corner from the rest of the classroom -small children's bench	Library corner: -books

Figure 3. Description of the school's rooms used for the experiments

Session	S1	S2	S3	S4	S5	S6	S7	S8	S9	S10
Room	R1	R1	R1	R1	R1	R3	- Child C : R3 - Other children : R4	- Child C: R3 - Other Children: R2	R2	R2

Figure 4. List of the school's room(s) used for each session

Sensor \ Behaviour Mode	1	2	3	4	5	6
Chin sensor	Wag the tail	Wag the tail	Emit "bark" sound while open-closing the mouth	Wag the tail	Open-close the mouth	Emit "bark" sound while open-closing the mouth
Head sensor	Open-close the mouth	Open-close the mouth	Move the head (head tilt oscillations)	Open-close the mouth	Wag the tail	Move the head (head tilt oscillations)
Back front sensor	Wag the tail	Wag the tail	Wag the tail	Walk forward, turn right, stand, turn left, walk backwards	Walk forward, turn right, stand, turn left, walk backwards	Walk forward, turn right, stand, turn left, walk backwards
Back middle sensor	Turn head (head pan oscillations)	Turn head (head pan oscillations)	Turn head (head pan oscillations)	Turn head (head pan oscillations)	Turn head (head pan oscillations)	Turn head (head pan oscillations)
Back rear sensor	Move the neck (oscillations)	Emit "bark" sound	Emit "drum" sound while wagging the tail	Emit "bark" sound and move the neck (oscillations)	Emit "drum" sound while wagging the tail	Emit "drum" sound while wagging the tail

Figure 5. Examples of Behaviour-Modes for the robot. Mapping between the external tactile sensors of the robot and its behaviours

Methodology of the approach. During the session, the child was invited to play with the Sony robotic pet Aibo. The experimenter took part in the experiment. The child was the major leader for play: the child was free to choose the game to focus on, the pace of play and he/she could engage in free-play (unconstrained play) with the robot and/or the experimenter; he/she was also free to engage in communication with the experimenter whenever he/she wanted. If the child appealed to the experimenter's participation, then the experimenter did take part in the game. If the child initiated verbal or non-verbal communication (e.g. smile, eye gaze) with the experimenter then the experimenter answered 'appropriately', e.g. (i) if the

child smiled to the experimenter, then the experimenter smiled back at the child, (ii) if the child looked at the experimenter the experimenter looked at the child (eye contact) and (iii) if the child initiated verbal communication, the experimenter answered appropriately, using words the child could understand, to faciliate social interaction. With respect to verbal communication, the experimenter tried to answer every question of the child and rewarded him/her verbally whenever appropriate (e.g. at the end of each play session, the experimenter told the child he/she played very well with Aibo and congratulated him/her.). Note that this approach is mainly child-centred, relies strongly on the child's capabilities of designing his/her own trajectory of progression and on total respect and consideration towards the child from the experimenter. In this sense, this approach draws inspiration from non-directive play therapy.

Beyond inspiration from non-directive play therapy, this approach adds a regulation process under specific circumstances which are detailed below:

a. *to prevent or discourage a repetitive behaviour*: If the child was starting or about to start a repetitive behaviour, the experimenter intervened and tried to help the child play a different game;[11]
b. *to help the child engage in play*: if the child did not engage in interaction with the robot, then the experimenter encouraged him/her to play with the robot, verbally and/or non-verbally (e.g. by stroking the robot and encouraging verbally imitation);
c. *to give a better pace to the game if already experienced by the child*: If the game was "standing still" but the child had already experienced it and had shown that he/she was capable of playing this specific game, then the experimenter could intervene straight away to confer a better pace to the game;
d. *to bootstrap a higher level of play*: if the child was about to reach a higher level of play but still needed some bootstrapping (some light guidance), the experimenter could provide it; note that the different levels of play are described in a Play Grid that is presented in the next subsection;
e. *to proactively ask questions related to affect or reasoning*: the experimenter could proactively ask the child simple questions related to affect or reasoning such as: "Do you think Aibo is happy today?" or "Do you like playing with Aibo?".

Note that (e) enables (i) testing the ability of the child to answer and/or (ii) showing the child a specific point for reasoning. We shall give several examples of levels of reasoning:

1. technical issue: show the child how to change the battery of the robot so that he/she can do it next time in a context of cooperative task;
2. ask the child if he/she thinks Aibo is happy;

3. help the child reason on causal effect: stimulation of a sensor implies a specific reaction of the robotic dog;
4. show the child that a reaction can be interpreted: e.g. if I press this specific button, then Aibo wags the tail; and wagging the tail can mean that Aibo is happy; thus if you press this button, you can show that Aibo is happy.

4.3.2 *Measures*

Each session was filmed unless the child explicitly asked not to be filmed which rarely happened. First, the experimenter viewed the video recordings and wrote down notes on the events constituting each session. These notes described the events in detail and contained as few interpretations as possible. As a second step, the experimenter analysed the data in terms of more abstract criteria that would enable her to identify, for each child, both the profile according to the three dimensions (Play, Reasoning and Affect) and the progress made over the 10 sessions. This methodology allows the researcher to first gather as much information as possible before deciding on the specific criteria; it has the advantage of not restricting the analysis to predefined criteria which might a posteriori turn out to be less optimal. This is especially relevant in the case of an exploratory study. This procedure follows the one described by Schatzman & Strauss (1973), stating that: "the researcher requires recording tactics that will provide him with an ongoing developmental dialogue" (p. 94). Schatzman & Strauss (1973) underline the importance of recording observations from the very beginning of research. They also suggest taking notes separately, categorizing notes into three different packages: (a) observational notes based on events, without interpretation; (b) theoretical notes representing an attempt to confer or denote the meaning from an observational note; (c) methodological notes dedicated to methodological comments.

Results of the experiments were analyzed according to three (intertwined) dimensions, respectively Play, Reasoning and Affect.

Play This study aims at testing the feasibility of this approach to encourage the child to learn new play skills and enable him/her to experience more and more complex play situations with respect to the following main criteria:

a. social aspect of play,
b. proportion of symbolic and/or pretend play,
c. understanding/use of causality,
d. ability to handle the pace of a specific play and possibly the chronology or the transitions between two logical segments of play.

That is why, concerning the dimension of play, what particularly matters is (1) to extract information qualitatively about play situations that the child has experienced

in each session, and (2) to see if the child really experienced a large repertoire of play and more complex levels of play gradually over the sessions.

For this purpose, a Play Grid was built (after the play sessions) based on the children's play observed during the experiments. This grid is exhaustive with respect to the variety of play situations which took place at least once during the experiments for at least one of the children. Besides, the different play situations were classified into 6 sets, each set denoting a specific level of complexity of play (Level 1 being the lowest and then gradually incrementing the level of complexity until Level 6). The level of complexity is defined according to four criteria:

a. social play,
b. proportion of pretend and/or symbolic play,
c. exploration of the use of causality/reaction,
d. chronology and/or number of different phases in the play, e.g. a simple reaction to a sensor is constituted of two phases while a search and rescue game involves many phases to handle chronologically: (i) initial situation, (ii) search phase, (iii) rescue phase, (iv) final situation.

The level of complexity is then deduced from an average evaluation over the four components which explains that the same level may contain play situations with a predominant component of '(d)' and others with a predominant component of '(b)'.[12] Consequently, within the same level of complexity, the different play situations are not ordered since they may be very different in nature. Ideally, the child would experience higher levels of play over time and, within the same level of complexity, play situations that are different in nature.

The systematic analysis with the grid for each child and each session shows the trajectory of each child (i.e. the profile of the child). Each cell in the grid is filled in if and only if it corresponds to a play situation experienced by the child at least once during that specific session; and the content depends on the play situation being acted proactively (i.e. child's own initiative) or reactively (i.e. the child was gently guided towards this play situation by the experimenter).

However, this grid is very enlightening with regard to children who manage to play socially and manage to diversify their play. For those who do not interact much with the robot and, when playing, tend to experience mainly solitary play through the exploration of the robot's features and behaviours, an additional tool to evaluate their progress was used. That evaluation was quantitative and relied on measuring for the whole duration of each session:

1. the total time spent in interaction with the robot,
2. the duration for each single uninterrupted phase (period) of pure interaction (note that the total duration is the sum of the duration of each single uninterrupted phase of play),

3. the amount of gestures imitated by the child and the number of gestures explicitly asked by the experimenter to be imitated.

Reasoning Through play, children can notably construct some understanding of social situations and gain experience of some situations they encountered while playing. If a child can reason about abstract concepts, infer mental states and make a sense of social rapport, it will be easier for him/her to play symbolically. Conversely, while the child experiences symbolic play, he/she manipulates abstract concepts such as inferring an emotion or handling social rapport. Both play styles and reasoning are therefore intertwined and both viewpoints should therefore be used to analyse the results of the experiments carried out for this study. Note that with respect to "Reasoning", what is particularly relevant is that both questions and answers emerg from play situations. The context of play enables the use of imagination, whereby Aibo may be assigned a specific role by the child, and it allows the child to attribute specific capacities to the robot such has having mental states (e.g. it enables him/her to imagine that Aibo is taking on a specific role and to make further assumptions on its mental state or its social status). Thus, the context of play enables the robotic pet to be attributed with mental states as well as a social role, and possibly moral standing. In this way, it is possible to explore the reasoning part of the coding manual developed by Kahn et al. (2003) for the analysis of children's conception of the Aibo robot, by exploring the four following categories used in Kahn et al. (2003): "Essence", "Mental States", "Social Rapport" and "Moral Standing". According to Kahn et al. (2003), those categories "reflect a 'quadrology' of children's conceptions of Aibo and Shanti".[13] For each of those four categories a list of related questions can be formulated (Kahn et al., 2003) that is provided in Fig 6.

Entity	Questions related
Essence	Does the child consider Aibo as an artifact or a biological entity?
Mental States	Does the child attribute mental states to Aibo? Does the child consider that the robot develops in terms of age for instance? Does the child consider Aibo has a personality? Does the child consider Aibo could live autonomously?
Social rapport	How does the child position Aibo relatively to himself/herself?
Moral standing	Can Aibo be physically or morally hurt? Can Aibo be held responsible for something? Can Aibo be punished when necessary? Could Aibo be praised?

Figure 6. Four categories proposed in Kahn et al. (2003) for the analysis of children's conceptions of Aibo. This table presents questions related to the four entities

Note that Kahn et al.'s coding manual has been developed in a different context: they targetted typically developing preschool children who only encountered Aibo once and afterwards immediately answered specific questions about "reasoning" (Kahn et al., 2003, 2006) – while answering questions, children could however carry on interacting with the robot. The context used in our study is different since the succession of sessions enabled the child to progressively build some reasoning and understanding, along with the progressive building of a shared space of expressions and routine activities between the child and the experimenter. Therefore, the reasoning related to the robot can be enriched. Besides, "reasoning" here is part of play in itself. In the study presented in this article, the context of play is actually used to enable the child to explore issues such as mental states or social rapport, and the robot in itself is a support for embodying such issues through the imaginary context that comes with play. Moreover, since the experimenter takes part in the experiments, not only social rapport between the child and the robot should be considered, but also the child's view on the notion of social rapport between the robot and the experimenter and between himself/herself and the experimenter. Consequently, here, the dimension of "Reasoning" is analysed as follows:

1. The main features of the four categories ("Essence", "Mental States", "Social Rapport" and "Moral Standing") are extracted from Kahn et al.'s coding manual (Kahn et al., 2003);
2. The issue of whether and how the child addresses those features is investigated for each child, in a perspective of questioning through play rather than giving firm answers.

Note that since the experimenter is not a therapist, and since the behaviour of children with autism might sometimes be interpreted differently from typically developing children, in the analysis we only consider events which are as much as possible objectively and reliably identifiable. Verbal events are particularly reliable events; they can be statements or questions arising from the child (major events) or answer to the experimenter's question (minor events). Below are some examples: (a) Essence: "He's a robot, he is a robot dog", "He has short teeth, he doesn't bite. Robot dogs don't bite, do some do?"; (b) Mental states: "Aibo is happy", "How old is Aibo", "Aibo, answer me, do you like toys?"; (c) Social Rapport: "It is your robot"; (d) Moral standing: the child accidently kicks the robot and apologizes verbally to the robot directly. Besides, in many cases, as already explained, reasoning and play are intertwined; for instance, when the child and the robot's relative social position in an enacted situation of pretend play is well-defined by the child (e.g. a competition with two participants, the child and Aibo), then the notion of social rapport is certainly addressed. Another example is a play situation of asking the robot about its mental states and answering with the activation of a sensor.

As a further step in reasoning, the child may tackle a more general issue related to his/her mental states for instance, or to social rapport, concerning himself/herself or even the experimenter. This is a relevant point for this study: it would show the potential reuse in another context of skills the child may develop or practise through reasoning about the robot during play.

Affect The 'Affect' dimension represents any expression indicating whether the child likes the robot or not, or if the child makes an assumption about the robot liking him/her. Here, only obvious signs (verbal expressions) of likes/dislikes are considered, (see Fig. 7 which provides the table of criteria for the coding of events related to affect). This is made in order to ensure that events considered as related to affect are clearly identifiable. For instance, a gentle stroke is not classified as an

1. Proactive (major) event related to affect:

(i) Child's statement or question referring directly to himself/herself liking the robot or the robot liking him/her. No hug or kiss from the child to the robot. Examples: "I like Aibo", "Aibo likes me".

(ii) Child's verbal compliment to/concerning the robot. No hug or kiss from the child to the robot. Examples: "good doggy", "nice dog", "he is a nice dog".

(iii) Child's hug to the robot, clearly identifiable, accompanied by a kind word from the child to/concerning the robot or verbal statement qualifying the hug. Example: the child hugs the dog and asks the experimenter to hug the dog: "Put your hands and hug, hug, hug!"

(iv) Child's kiss to the robot, clearly identifiable, accompanied by a kind word from the child to/concerning the robot. Example: the child gives a kiss to Aibo after saying "Goodbye Aibo, have a good sleep"

2. Reactive (minor) event related to affect:

(i) Child's answer to a question about himself/herself liking the robot or the robot liking the child. Example: the experimenter asks the child: "Is it a nice robot?" and the child answers "Yes".

(ii) Child's answer to a question about himself/herself being happy to play with the robot. Example: the experimenter asks the child: "Are you happy playing with the robot?" and the child answers "Yes".

Note, reactive events related to affect are considered very cautiously in this study; they are not considered as sufficient to make firm deductions about the child addressing the notion of "Affect".

Figure 7. Criteria for coding events related to Affect. An event is related to 'Affect' if it corresponds to one of the items provided in the table; in some of the following figures, events related to affect are qualified by a corresponding code: the code of an event related to affect is given by its corresponding item's index, e.g. "I like Aibo" is [1i]

event related to affect in this study, neither a gesture such as a kiss or a hug, if it is not accompanied by an appropriate child's statement.

4.4 Coding and reliability

Inter-rater reliability testing was carried out for each of the three dimensions: play, reasoning and affect. A second coder who was not familiar with the aims of the study re-coded part[14] of the data. Good reliability was shown: (a) on play, 80.75% agreement (13*min*50*s* of videos coded divided among two children, Child E and Child C); (b) on reasoning, 80.35% agreement (18*min*24*s* of videos coded divided among two children, Child E and Child F); (c) on affect, 93.35% agreement (22*min* of Child C's videos coded).

5. Results

In the following we provide case study evaluations for each child.

Child A Child A showed some apprehension towards the robot and did not interact at all during the five first sessions. The experimenter therefore decided not to require the child to come for the following sessions and let the child proactively decide whether he wanted to take part in the further trials or not. In the last session (Session 10), Child A proactively came for the trial. In that session he engaged in an interaction with the robot with the help of the experimenter: one interaction event happened between the child and the robot, during which the experimenter showed the child how to stroke the robot and the child imitated (Fig. 8). Afterwards, the child showed both signs of light apprehension (he moved his body slightly backwards) and enjoyment (he smiled).

Child B Child B took part in 9 sessions (Fig. 9). Child B naturally showed attempts to play with the laptop rather than with the robot. It was a big challenge to get the child away from the laptop and get his attention focused on something else. The experimenter used a simple trick by hiding the laptop with a cloth. But for practicality reasons (e.g. to connect or reconnect Aibo during the session), the cloth had to be removed from time to time during the session thus introducing an important source of distraction for Child B. Progressively, the child seemed to have understood that he was allowed to occasionally have a look at the laptop (as part of his well-being) but that he should mostly engage in interactions with the robot. The table provided in Fig. 10 shows the average amount of time Child B spent engaging in play with the robot during each session. The tendency is clearly that the child

		1	2	3	4	5	6	7	8	9	10
L	*Solitary* Exploration										
1	"Imitation" of robot's bark										
	Solitary mirror play – look at oneself in the robot's reflecting face										
L 2	*"Pre-social" or basic-social* exploration – stroke Aibo immediately after the experimenter (possibly basic imitation of the gesture)										P
L	*Social* exploration (social play)										
3	Simple Bite/Save or Give/Food - no use of the sensors										
	Position or locomotion game – with verbal qualification of the game										
	Cooperative technical task: change the battery, or turn on/off Aibo										
	Verbal order towards Aibo: e.g. "sit", "walk", "wake up"										
	Basic pretend & social play – imitate Aibo's snoring & verbal comment										
	Basic play on affective gestures – give/receive a kiss and/or a lip to/from Aibo										
	Repeat after me - ask the experimenter to repeat verbal expressions										
	Look at Aibo through the camera (Possibly stroke Aibo & look at its reaction through the camera)										
	Speak French with Aibo - e.g. "Hello" or "Bye-Bye" in French										
	Show Aibo to other children (social play) Express verbally the willing/intention to show Aibo to the other children										
	Simple play with accessory (symbolic play)										
	Social Mirror play (social play) - look at oneself (and possibly at the experimenter) in the robot's reflecting face & express verbal comments, e.g. "Look at my arm!"										
	Social Hug – hug Aibo & ask the experimenter or the second researcher to hug Aibo										
L 4	*Complex* Give Food/Drink (cause-reaction play & symbolic play & social play) - use of sensors										
	Complex Bite/Save (cause-reaction play & pretend play & cooperative play) - use of sensors										
	Complex turn off Aibo to sleep (symbolic play)										
	Speak directly to Aibo about Aibo's feeling (symbolic play)										
	Cause-reaction play & mental states: Ask a question to Aibo (e.g. identity, feeling), answer with a sensor										
	Cause-reaction play, Aim at a physical reaction of the robot, show it with a sensor										
	Cause-reaction play & basic pretend play, "caught on the act"										
	Telling a story										
L 5	Cause-reaction play and explicit Social rapport: Ask a question to Aibo, answer with a sensor (e.g. press the sensor which opens the mouth), translate verbally the answer for the experimenter										
	Symbolic & pretend play Complex play with an accessory										
	Symbolic & pretend play Complex nap with Aibo										
	Symbolic & extrapolation play: "RobotCat" - Speak about the idea of a robotic cat (possibly imagine how one would play with it)										
	Causal composition of plays: Bite/Save & Give Food/Drink										
	Causal composition of plays: Kiss & Bite/Save										
	Pretend play & causal reaction & social rapports: Ask verbally Aibo to act a situation, use of sensors										
L 6	Pretend play & focus on Aibo's mental states: Mimic Aibo's cry, and explain Aibo is never crying but pretending to cry										
	Pretend play & social rapports: Look after Aibo and set up rules										
	Pretend & symbolic & chronological play & social rapports: Search and rescue										
	Pretend & symbolic play & social rapport & cause-reaction play & chronological play: competition (drink fast) between the child or the experimenter and Aibo ; the non-competitor activates Aibo's sensor										

played longer with the robot in the last two sessions than in the previous ones and almost doubled his play time between the 9th and 10th session. If we consider in detail the duration of single phases of play, i.e. uninterrupted periods of time when the child continuously played with the robot, then, again, this table shows that the child experienced longer uninterrupted periods of play with the robot during the last sessions. Typically, two uninterrupted periods of play are often separated by an attempt of the child to play with the laptop. This shows that the child progressively learnt to focus more and more on the robot and on engaging in play with the robot. Nevertheless, the experimenter also often intervened to help the child carry on playing and keep focusing his total attention to the robot; this intervention usually happened in two ways: (a) encouraging and rewarding the child verbally, or (b) showing an example, e.g. stroking the robot and asking for the child to do the same. In this context, '(b)' is very relevant indeed since the child does not speak verbally and encouraging imitation is favourable for both relaunching the child's engagement in play and bootstrapping social play. It should be noted that in this specific context, imitation is very rudimentary: the experimenter either touches a specific sensor or gently strokes the robot (e.g. on the head) and explicitly asks the child to do the same. The child is considered to imitate the experimenter's gesture if he exhibits the same type of gesture within 10 seconds, i.e. either by touching a sensor or stroking, and if the gesture is applied on the same part of the robot's body; for instance, (i) the experimenter touches the head sensor and, within 10 seconds, the child presses the same sensor (with or without activation depending on the child's precision of touch) ; or (ii) the experimenter gives a gentle stroke on the back of the robot and, within ten seconds, the child gives a stroke on the back of the robot. Results show that Child B progressively experienced more situations of imitation. Besides, they also reveal that during the last session he imitated some gestures proactively, i.e. without being explicitly asked by the experimenter to imitate.

Figure 8. Child A. Play Grid. The first column describes the corresponding level of play, the second column details the various play situations for each level that the child experienced at least once; the following columns refer to the sessions, ordered chronologically. The table is then completed according to the following rules: (a) if the child did not experience the play situation during the specific session, leave the corresponding cell blank; (b) if the child experienced the specific play situation at least once during the session, then write "P" (if the child experienced it proactively only – i.e. it was his/her own initiative). Write "r" if the child never experienced it proactively (only reactively: the experimenter guided the child towards the play situation). Write "B" if the child experienced this play situation several times, sometimes proactively and sometimes reactively. Note that Child A did not take part in the play sessions 6, 7, 8 and 9

		1	2	3	4	5	6	7	8	9	10
L 1	*Solitary* Exploration	P	B	B	P	r			B	P	B
	"Imitation" of robot's bark										
	Solitary mirror play – look at oneself in the robot's reflecting face										
L 2	*"Pre-social" or basic-social* exploration – stroke Aibo immediately after the experimenter (possibly basic imitation of the gesture)								r	r	B
L 3	*Social* exploration (social play)										
	Simple Bite/Save or Give/Food - no use of the sensors										
	Position or locomotion game – with verbal qualification of the game										
	Cooperative technical task: change the battery, or turn on/off Aibo										
	Verbal order towards Aibo: e.g. "sit", "walk", "wake up"										
	Basic pretend & social play – imitate Aibo's snoring & verbal comment										
	Basic play on affective gestures – give/receive a kiss and/or a lip to/from Aibo										
	Repeat after me - ask the experimenter to repeat verbal expressions										
	Look at Aibo through the camera (Possibly stroke Aibo & look at its reaction through the camera)										
	Speak French with Aibo - e.g. "Hello" or "Bye-Bye" in French										
	Show Aibo to other children (social play) Express verbally the willing/intention to show Aibo to the other children										
	Simple play with accessory (symbolic play)										
	Social Mirror play (social play) - look at oneself (and possibly at the experimenter) in the robot's reflecting face & express verbal comments, e.g. "Look at my arm!"										
	Social Hug – hug Aibo & ask the experimenter or the second researcher to hug Aibo										
L 4	*Complex* Give Food/Drink (cause-reaction play & symbolic play & social play) - use of sensors										
	Complex Bite/Save (cause-reaction play & pretend play & cooperative play) - use of sensors										
	Complex turn off Aibo to sleep (symbolic play)										
	Speak directly to Aibo about Aibo's feeling (symbolic play)										
	Cause-reaction play & mental states: Ask a question to Aibo (e.g. identity, feeling), answer with a sensor										
	Cause-reaction play, Aim at a physical reaction of the robot, show it with a sensor										
	Cause-reaction play & basic pretend play, "caught on the act"										
	Telling a story										
L 5	Cause-reaction play and explicit Social rapport: Ask a question to Aibo, answer with a sensor (e.g. press the sensor which opens the mouth), translate verbally the answer for the experimenter										
	Symbolic & pretend play Complex play with an accessory										
	Symbolic & pretend play Complex nap with Aibo										
	Symbolic & extrapolation play: "RobotCat" - Speak about the idea of a robotic cat (possibly imagine how one would play with it)										
	Causal composition of plays: Bite/Save & Give Food/Drink										
	Causal composition of plays: Kiss & Bite/Save										
	Pretend play & causal reaction & social rapports: Ask verbally Aibo to act a situation, use of sensors										
L 6	Pretend play & focus on Aibo's mental states: Mimic Aibo's cry, and explain Aibo is never crying but pretending to cry										
	Pretend play & social rapports: Look after Aibo and set up rules										
	Pretend & symbolic & chronological play & social rapports: Search and rescue										
	Pretend & symbolic play & social rapport & cause-reaction play & chronological play: competition (drink fast) between the child or the experimenter and Aibo ; the non-competitor activates Aibo's sensor										

Figure 9. Child B. Play Grid. See Fig. 8 for a detailed caption. Note that Child B was away for Session 7

	Total duration of play (min: sec)	Repartition of the play time in single phases of play (min:sec and + between 2 single phases)	Aspects of imitation: In each single phase of play, numbers of gestures:		Verbal expression involving either the word 'dog' or 'robot'
			Imitated by the child	Explicitely asked by the experimenter to be imitated	
Session1	0:06	0:06	0	0	
Session2	1:30	1:00 + 0:30 (mostly looking attentively at Aibo)	0	0	
Session3	0:40	0:40	0	0	
Session4	Almost null	Almost null	0	0	'The little dog was easy'
Session5	0:15	0:15 the experimenter helps by holding the child's hand to show him	0	0	
Session6	0:00	0:00	0	0	
Session7	away				
Session8	1:05	1:05	1	2	
Session9	2:21	0:40 +1:16 +0:16	0 +1 +0	0 +2 +0	
Session10	5:24	0:20 +1:47 +0:18 +2:46	0 +3 +0 +3	0 +3 +0 +1	

Figure 10. Child B. Dimension of play: quantitative results. For each session, the following indicators are reported: (a) total duration of play; (b) duration for each specific single session of play ; (c) aspects of imitation with respect to (i) the occurrence of gestures (touch or stroke of the robot) that the child imitated and (ii) the occurrence of gestures that the experimenter explicitly asked the child to imitate; (d) verbal expressions including the word "dog" or "robot"

Concerning the "Reasoning" dimension, Child B did not address the issue verbally. Thus, no firm conclusions should be drawn. However, the detailed study of the child's gestures shows that the exploration of the child became progressively richer over the sessions. The child varied his position relative to the robot, from sitting to kneeling and lying, and thus looked at the robot from various viewpoints. Moreover, he progressively varied his way of touching the robot: during the first sessions, he progressively abandoned random-like touch to develop more targeted touch. Note that targeted touch can be, for instance, trying to touch a single sensor precisely or stroke the robot gently and then activate many sensors. Besides, during the last session, the child experienced proactively a combination of two previous sensor activations: first, he imitated the experimenter and stroked the back of the robot; then he imitated the experimenter again and touched the head; third, he simultaneously activated the robot's back and head sensors.

Concerning the third dimension, "Affect", no event that was related to affect (with respect to Fig. 7) was recorded.

Child D Child D was away for Session 3 and Session 6 and therefore took part in 8 sessions in total. The analysis of the Play Grid in Fig. 11 shows that Child D played mostly solitarily. He engaged largely in exploratory play which became progressively more and more enriched. Two main aspects objectively illustrate the phenomenon (a) a progressive change of position (from sitting orthogonal to the robot and not facing the experimenter to facing the robot and the experimenter) and (b) a more diversified way of touching the sensors. Moreover, the child practised "solitary mirror play" frequently. It consists of looking at one's own image in the robot's reflecting face. Child D experienced situations of looking at his image with other reflecting surfaces too, such as a window partially reflecting, or a mirror perfectly reflecting (room R2 contained a mirror). All of these play situations, consisting of looking at one's own image, were often fascinating for Child D, and sometimes prevented him from engaging in other kinds of play situations. Besides, Child D did not experience play involving explicitly causal reactions, such as showing a specific reaction of the robot through the sensors' activation.

However, progressively, Child D experienced situations with some components of social play. From a cooperative point of view, the child did take part, both reactively and proactively in cooperative technical tasks such as turning on the robot. Furthermore, Child D, who mostly speaks by using onomatopeia did develop some ways of expressing himself, by dancing in front of the mirror and/or the robot and even probably telling a story by not using proper words but onomatopeia. The situation described below, that Child D experienced, may actually be interpreted, with caution, as a storytelling situation: Child D chronologically (a) pressed the button to "wake up" Aibo (i.e. turn Aibo on), then

		1	2	3	4	5	6	7	8	9	10
L 1	*Solitary* Exploration	P	P		P			P	P	P	P
	"Imitation" of robot's bark										
	Solitary mirror play – look at oneself in the robot's reflecting face	P			P	P		P	P	P	P
L 2	*"Pre-social" or basic-social* exploration – stroke Aibo immediately after the experimenter (possibly basic imitation of the gesture)				P	P		P	B		
L 3	*Social* exploration (social play)										
	Simple Bite/Save or Give/Food - no use of the sensors										
	Position or locomotion game – with verbal qualification of the game										
	Cooperative technical task: change the battery, or turn on/off Aibo				P	P		B	P	P	B
	Verbal order towards Aibo: e.g. "sit", "walk", "wake up"										
	Basic pretend & social play – imitate Aibo's snoring & verbal comment										
	Basic play on affective gestures – give/receive a kiss and/or a lip to/from Aibo										
	Repeat after me - ask the experimenter to repeat verbal expressions										
	Look at Aibo through the camera (Possibly stroke Aibo & look at its reaction through the camera)										
	Speak French with Aibo - e.g. "Hello" or "Bye-Bye" in French										
	Show Aibo to other children (social play) Express verbally the willing/intention to show Aibo to the other children										
	Simple play with accessory (symbolic play)										
	Social Mirror play (social play) - look at oneself (and possibly at the experimenter) in the robot's reflecting face & express verbal comments, e.g. "Look at my arm!"										
	Social Hug – hug Aibo & ask the experimenter or the second researcher to hug Aibo										
L 4	*Complex* Give Food/Drink (cause-reaction play & symbolic play & social play) - use of sensors										
	Complex Bite/Save (cause-reaction play & pretend play & cooperative play) - use of sensors										
	Complex turn off Aibo to sleep (symbolic play)										
	Speak directly to Aibo about Aibo's feeling (symbolic play)										
	Cause-reaction play & mental states: Ask a question to Aibo (e.g. identity, feeling), answer with a sensor										
	Cause-reaction play, Aim at a physical reaction of the robot, show it with a sensor										
	Cause-reaction play & basic pretend play, "caught on the act"										
	Telling a story									P	P
L 5	Cause-reaction play and explicit Social rapport: Ask a question to Aibo, answer with a sensor (e.g. press the sensor which opens the mouth), translate verbally the answer for the experimenter										
	Symbolic & pretend play Complex play with an accessory										
	Symbolic & pretend play Complex nap with Aibo										
	Symbolic & extrapolation play: "RobotCat" - Speak about the idea of a robotic cat (possibly imagine how one would play with it)										
	Causal composition of plays: Bite/Save & Give Food/Drink										
	Causal composition of plays: Kiss & Bite/Save										
	Pretend play & causal reaction & social rapports: Ask verbally Aibo to act a situation, use of sensors										
L 6	Pretend play & focus on Aibo's mental states: Mimic Aibo's cry, and explain Aibo is never crying but pretending to cry										
	Pretend play & social rapports: Look after Aibo and set up rules										
	Pretend & symbolic & chronological play & social rapports: Search and rescue										
	Pretend & symbolic play & social rapport & cause-reaction play & chronological play: competition (drink fast) between the child or the experimenter and Aibo ; the non-competitor activates Aibo's sensor										

(b) stood in front of the wall mirror in the room, still watching Aibo "waking up"; (c) once Aibo had "woken up", the child started dancing and using onomatopeia in front of the mirror. At some point, the robot disconnected. During the whole process the experimenter told Child D many times that she thought he was telling a story and asked him if she was right. She got no answer. When the robot disconnected the child stopped dancing and the experimenter reiterated her question: "Was it a story that you were telling me? Yes or no?"and the child answered "Yes". Then she asked: "Can you tell me another story, yes or no?" and the child answered "yes". Then the child repeated the same succession of behaviours '(a)', '(b)' and '(c)' and she asked: "Is the story about a boy?" And he answered "Yes". It is worthy of note here that the child might have simply repeated the word 'yes' after each question without giving a 'real' answer to the questions. Nonetheless, that example shows how the child may have progressively opened up to more communication with his surrounding social environment for play (notably the experimenter).

This storytelling situation took place in the last sessions while the child was starting to answer some questions about reasoning as well as using proactively verbal expressions to express intention. An in depth study of the verbal answers the child gave shows that over the first sessions, the child almost only answered "yes" or "no", whenever he answered. Then, progressively, the child answered some questions by repeating words from the question: e.g. in Session 4 the experimenter asked "Do you want to play with the robot or go back to the classroom?". The child answered: "play with the robot". And in the last two sessions, the child did use expressions to express his own intentions; for instance, the expression "sitting down" means that he wants to remain sitting down on the ground to carry on playing with the robot. In Session 9, the experimenter actually asked the child: "Do

Figure 11. Child D. Play Grid. The first column describes the corresponding level of play, the second column details the various play situations for each level that the child experienced at least once; the following columns refer to the sessions, ordered chronologically. The table is then completed according to the following rules: (a) if the child did not experience the play situation during the specific session, leave the corresponding cell blank; (b) if the child experienced the specific play situation at least once during the session, then write "P" (if the child experienced it proactively only – i.e. it was his/her own initiative). Write "r" if the child never experienced it proactively (only reactively: the experimenter guided the child towards the play situation). Write "B" if the child experienced this play situation several times , sometimes proactively and sometimes reactively. Note that Child D was away for Session 3 and Session 6

you want to go back to the classroom or play with him (the robot)?" and the child answered "play with him". Then later in the session, the experimenter asked the question "Shall we go back to the classroom now?" and the child answered: "Sitting down". During the last session, the child reused exactly the same expression ("sitting down") to answer the experimenter's question: "Would you like to go back to the classroom soon?".

Regarding the analysis of the reasoning dimension, the child answered reactively very basic questions about Aibo's mental states, such as "Do you think Aibo is happy today?" or about his own mental state: "Do you like playing with the robot?" but there was no proactivity from the child with respect to mental states.

Concerning "Social rapport", the child progressively grasped the fact that Aibo belonged to the experimenter. In the first sessions, the experimenter had to explain many times to the child that he could not take the robot with him back to the classroom. In contrast, at the end of the last session, the child hesitated a short time and gave the robot back to the experimenter proactively. Apart from that, the child did not explicitly show any reasoning on "Social rapport" or on Aibo's "Moral standing".

The dimension of Affect has been mostly addressed indirectly (Fig. 12), through simple questions from the experimenter: in Session 4, the child answered affirmatively to the following questions: (a) "Is it a nice robot?" and (b) "Are you happy playing with the robot?". Later, in Session 9, the child answered affirmatively to the question "Do you think Aibo likes you?" And in Session 10, the child answered affirmatively to the question "You like the robot?". Note that since these inputs did not emerge proactively we should be careful with too much interpretation. Nonetheless, it should be underlined that most of the time the child said he preferred playing with the robot rather than going back to the classroom, which shows the child was having fun playing with the robot. Note, the experimenter is aware that the child may just have given a stereotypical answer.[15]

Child C Child C was away for Session 7 and thus took part in 9 sessions in total (note that in Session 6 she had a very limited time of play, approximately 10 minutes, because of a class trip). The Play Grid in Fig. 14 shows that Child C experienced more and more complex levels of play during the sessions (see Fig. 13). She experienced play situations involving the activation of a specific sensor to generate a precise reaction only a bit. She rather proactively experienced firstly play situations where "affect" is largely addressed (e.g. "Social Hug"). Secondly, she developed play situations where the robot embodied a character in a story she was telling. Finally, in a third and last phase, she initiated play situations where she was

Session	**Events objectively related to Affect** (ordered chronologically with respect to first appearance, event only mentioned once per session)
S1	
S2	· [2i] **“Do you like it?”** (Experimenter); **“Yes”** (Child D)
S3	
S4	· [2i] **“Is it a nice robot?** (Experimenter); **“Yes”** (Child D); ·[2ii] **“You are happy playing with the robot?** (Experimenter); **“Yes”** (Child D)
S5	
S6	
S7	
S8	
S9	· [2i] **“Do you think Aibo likes you?”** (Experimenter); **“Yes”** (Child D)
S10	· [2i] **“You like the robot?”** (Experimenter); **“Yes”** (Child D)

Figure 12. Child D. Events related to Affect. Events are separated by bullet points, and provided with their context (normal font) in the table. Events written in bold are coded according to Fig. 7 (the code is provided in brackets in front of the event); please note that when the child answers a question, the event in itself is the child's answer, but, in this table, in order to make it clear to the reader, the question that the answers refers to is also written in bold

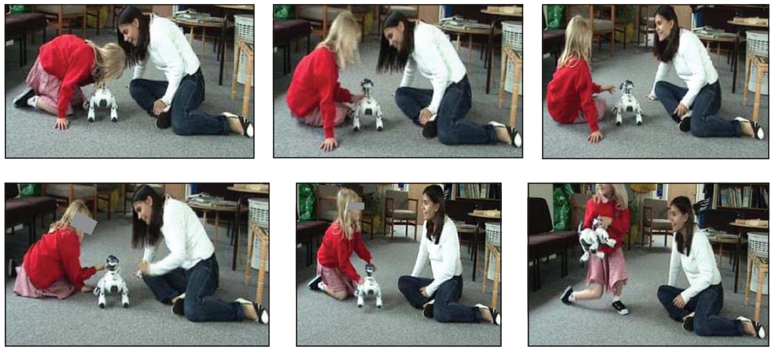

Figure 13. Child C involved in social play with the experimenter. Two sequences are displayed, one on each line. Each sequence is organised chronologically; on the first line, picture on the right and on the second line, picture in the middle, Child C is making eye contact with the experimenter

able to tackle issues on social rapport or mental states (Session 10: "looking after Aibo and set up rules" and "search and rescue" play situations).

The "looking after Aibo" game dealt with deciding that she and the experimenter would take care of Aibo, and Child C proactively suggested that, as a

		1	2	3	4	5	6	7	8	9	10
L 1	*Solitary* Exploration										
	"Imitation" of robot's bark				P	P			P		
	Solitary mirror play – look at oneself in the robot's reflecting face										
L 2	*"Pre-social" or basic-social* exploration – stroke Aibo immediately after the experimenter (possibly basic imitation of the gesture)										
L 3	*Social* exploration (social play)	P	P	P	P	P	P		P	P	P
	Simple Bite/Save or Give/Food - no use of the sensors						r				P
	Position or locomotion game – with verbal qualification of the game	P				P	P		P		
	Cooperative technical task: change the battery, or turn on/off Aibo		P	P	P		r		r	P	
	Verbal order towards Aibo: e.g. "sit", "walk", "wake up"		P	P	P				P	P	P
	Basic pretend & social play – imitate Aibo's snoring & verbal comment		P								
	Basic play on affective gestures – give/receive a kiss and/or a lip to/from Aibo			P	P	P	P				
	Repeat after me - ask the experimenter to repeat verbal expressions										P
	Look at Aibo through the camera (Possibly stroke Aibo & look at its reaction through the camera)				P						
	Speak French with Aibo - e.g. "Hello" or "Bye-Bye" in French										
	Show Aibo to other children (social play) Express verbally the willing/intention to show Aibo to the other children										
	Simple play with accessory (symbolic play)										
	Social Mirror play (social play) - look at oneself (and possibly at the experimenter) in the robot's reflecting face & express verbal comments, e.g. "Look at my arm!"										
	Social Hug – hug Aibo & ask the experimenter or the second researcher to hug Aibo			P							
L 4	*Complex* Give Food/Drink (cause-reaction play & symbolic play & social play) - use of sensors								B	B	P
	Complex Bite/Save (cause-reaction play & pretend play & cooperative play) - use of sensors										
	Complex turn off Aibo to sleep (symbolic play)										
	Speak directly to Aibo about Aibo's feeling (symbolic play)										
	Cause-reaction play & mental states: Ask a question to Aibo (e.g. identity, feeling), answer with a sensor						P				
	Cause-reaction play, Aim at a physical reaction of the robot, show it with a sensor						r		P		r
	Cause-reaction play & basic pretend play, "caught on the act"										
	Telling a story				P		P		P	P	
L 5	Cause-reaction play and explicit Social rapport: Ask a question to Aibo, answer with a sensor (e.g. press the sensor which opens the mouth), translate verbally the answer for the experimenter										
	Symbolic & pretend play Complex play with an accessory										
	Symbolic & pretend play Complex nap with Aibo										
	Symbolic & extrapolation play : "RobotCat" - Speak about the idea of a robotic cat (possibly imagine how one would play with it)										
	Causal composition of plays: Bite/Save & Give Food/Drink										
	Causal composition of plays: Kiss & Bite/Save										
	Pretend play & causal reaction & social rapports: Ask verbally Aibo to act a situation, use of sensors										
L 6	Pretend play & focus on Aibo's mental states: Mimic Aibo's cry, and explain Aibo is never crying but pretending to cry										
	Pretend play & social rapports: Look after Aibo and set up rules										P
	Pretend & symbolic & chronological play & social rapports: Search and rescue										P
	Pretend & symbolic play & social rapport & cause-reaction play & chronological play: competition (drink fast) between the child or the experimenter and Aibo ; the non-competitor activates Aibo's sensor										

Figure 14. Child C. Play Grid. See Fig. 11 for a detailed caption. Note that Child C was away for Session 7

consequence, she and the experimenter would have to define rules the robot would have to respect; and she enumerated the rules (among them, a detailed list of what the robot is not allowed to eat, and the statement: "dogs must go outside and must walk", followed by "I need to make him walk"). This game also gave rise to proactive inferences of state, the child even saying: "Look! He is smiling!" in the proper context. The social status that she took of taking care of Aibo led her to show the experimenter how to do specific things such as to make Aibo go forward: "You see, you must do like this, see".

Furthermore, this game was followed by a "search and rescue game" which was extremely rich in many ways:

a. The child led the rhythm, the pace, and the three steps of the play situation (chronologically):
   - step 1: initial situation where Aibo is lost, the goal of finding Aibo is stated,
   - step 2: the experimenter and the child are looking for the dog,
   - step 3: final situation: the experimenter and the child find the dog.
b. The child slightly extended step 2 over time so that she could deal with emotional states, particularly sadness: "You think we've lost him forever" said Child C; "Oh, that's sad" said the experimenter; and the child replied: "I think we're sad actually" thus conferring a socio-dramatic dimension to the current play situation.
c. During step 3, when the robot was found, the child introduced some reasoning about categories: she introduced the notion that it might be a robot other than Aibo that she and the experimenter had found; she introduced this reasoning step by step and she might not have been really at ease with these concepts, but the point is that she practised them through experiencing them: Child C's reasoning started with "Oh no, there are two Aibos here" and, after several steps in the reasoning, she drew the following conclusion: "No there are two dogs, only one Aibo. The clever one!" and she threw up her hands accompanied by a big smile. Again, what is illustrated here is that both "reasoning" and "play" dimensions are highly intertwined.

Concerning the notion of "Essence" for the Reasoning dimension, Child C mixed the use of artifacts and biological statements such as saying within the same session: "He's a robot, he's a robot dog" and "Nice dog", "He is a nice dog", "I love dogs", "A boy or a girl?" (Session 10).

Except in the last session, the notion of "Mental states", was addressed mostly reactively: the child answered questions asked by the experimenter such as "Do you think Aibo is hungry?" (which usually initiates the game "Give food/drink"). There were two exceptions: (a) the child proactively said that the robot liked her, and (b) the child could sometimes refer to mental states when telling stories she adapted from well-known children's books. During the last session, the child proactively referred to mental states of the robot as mentioned above in both "look after" and "search and rescue" play situations. During the "look after" play situation, she said: "We play, want to make the dog happy, make the dog feel pretty".

Moreover, as already mentioned above too, she experienced "Social rapport" a lot e.g. either simply by saying (in Session 9) "Look at Aibo, Aibo is your dog" or in taking on specific social roles in more elaborated play situations (e.g. in Session 10, during "look after" and "search and rescue" games).

Concerning "Moral standing", no objective event related to it happened.

The dimension of "Affect" played an important role for the child (Fig. 16). In Session 1 already, she started saying "good doggy" with respect to the robot. Then, in Session 3 she introduced the notion of social hug (see Fig. 15), which consisted in asking the experimenter (or the second researcher present) to help her hug the dog: "Put your hands and hug, hug, hug" Child C asked. Later in the same session, as well as in Session 4, the child said, "The dog really likes me". Note that end of Session 3 is the first time she answered the question "Do you like it (Aibo)?" (she answered affirmatively). From that session onwards, the child confirmed several times the fact that Aibo liked her (e.g. Session 4 "The dog really likes me") and that she liked Aibo (e.g. in Session 10: "I love Aibo" and "Nice dog").

Figure 15. Child C's social hug to the robot. Photos ordered chronologically. The child brings the robot to a second researcher (who helped out during this trial) while saying "Put your hands and hug, hug, hug" and both of them hug the dog. In the third picture from the left, Child C makes eye contact with the researcher

Session	**Events objectively related to Affect** (ordered chronologically with respect to first appearance, event only mentioned once per session)
S1	· [1ii] **"Good doggy"** (Child C) while stroking the robot and looking at the experimenter (eye contact)
S2	
S3	· [1iii] **"Help me hug the dog: put your hands and hug, hug, hug"** (Child C) while bringing the robot near the assistant and showing how to hug · [1ii] **"Good doggy"** (Child C) · [1i] **"The dog really likes me"** (Child C). The experimenter answer "yes" · [2i] **"Do you like it?** (Experimenter). **"Yes"** (Child C)
S4	· [1ii] **"Good doggy"** (Child C), while stroking the robot · [1i] **"The dog really likes me"** (Child C) and she starts mimicking the noise that would do the dog by lapping her.
S5	· [1ii] **"Good doggy"** (Child C) and she looks at the experimenter; "yes very good doggy" (Experimenter).
S6	
S7	
S8	· [1ii] **"Good doggy"** (Child C) after the robot has "woken up" (i.e. is connected)
S9	· [2i] **Are you happy to see Aibo?** (Experimenter); **"Yes"** (Child C)
S10	· [1ii] **"Nice dog"** (Child C) · [1i] **"I love Aibo. I love Aibo"** (Child C) and she strokes the robot · [1ii] **"Good boy, good boy"** (Child C) and she strokes the robot · [1i] "Do you like the walk C, please tell me? (Experimenter); **"Yes, this is all about dogs like me"** (Child C) · [2i] **You like Aibo, right?** (Experimenter); **"Yes"** (Child C)

Figure 16. Child C. Events related to Affect. See caption of Fig. 12 for details

Child E. Child E took part in the 10 sessions of experiments. The Play Grid in Fig. 17 shows that Child E progressively experienced more and more complex levels of play over the sessions. During the first sessions, he attentively explored the reactions of the robot and in the following sessions, he experienced more and more simple causal reactions through the following games: (a) "ask about a feeling, answer with a sensor", e.g. in Session 10 the child asked: "are you happy?" and pressed the head button which made the robot wave the mouth as to say "yes". (b) "aim at a physical reaction, show it with sensors": e.g. the experimenter asked "Do you think Tornado (the name the child gave to the robot) can wag the tail today?" and Child E activated the right sensor at the first attempt and commented: "That's the tail one". Child E also proactively played the game of giving food or drink to the robot as well as a cooperative play situation of Bite/Save (see Fig. 18). Bite/Save play situation consisted of two chronological steps: (i) the robot bit the finger of either the child

		1	2	3	4	5	6	7	8	9	10
L 1	*Solitary* Exploration										
	"Imitation" of robot's bark										
	Solitary mirror play – look at oneself in the robot's reflecting face			P							
L 2	*"Pre-social" or basic-social* exploration – stroke Aibo immediately after the experimenter (possibly basic imitation of the gesture)										
L 3	*Social* exploration (social play)	P	P	P	P	P	P	P	P	P	P
	Simple Bite/Save or Give/Food - no use of the sensors						r	r			
	Position or locomotion game – with verbal qualification of the game	P		P					P		P
	Cooperative technical task: change the battery, or turn on/off Aibo	P	P	P	P	B	P	P	P	P	P
	Verbal order towards Aibo: e.g. "sit", "walk", "wake up"		P			P					
	Basic pretend & social play – imitate Aibo's snoring & verbal comment										
	Basic play on affective gestures – give/receive a kiss and/or a lip to/from Aibo										
	Repeat after me - ask the experimenter to repeat verbal expressions										
	Look at Aibo through the camera (Possibly stroke Aibo & look at its reaction through the camera)										
	Speak French with Aibo - e.g. "Hello" or "Bye-Bye" in French										
	Show Aibo to other children (social play) Express verbally the willing/intention to show Aibo to the other children										
	Simple play with accessory (symbolic play)							P			
	Social Mirror play (social play) - look at oneself (and possibly at the experimenter) in the robot's reflecting face & express verbal comments, e.g. "Look at my arm!"		P	P		P	P	P	P		
	Social Hug – hug Aibo & ask the experimenter or the second researcher to hug Aibo										
L 4	*Complex* Give Food/Drink (cause-reaction play & symbolic play & social play) - use of sensors						B	B	B	B	B
	Complex Bite/Save (cause-reaction play & pretend play & cooperative play) - use of sensors			P		B	r	P	P	P	P
	Complex turn off Aibo to sleep (symbolic play)							P			
	Speak directly to Aibo about Aibo's feeling (symbolic play)						P	P		P	P
	Cause-reaction play & mental states: Ask a question to Aibo (e.g. identity, feeling), answer with a sensor					r		r		P	P
	Cause-reaction play, Aim at a physical reaction of the robot, show it with a sensor			r		B			r		r
	Cause-reaction play & basic pretend play, "caught on the act"										P
	Telling a story										
L 5	Cause-reaction play and explicit Social rapport: Ask a question to Aibo, answer with a sensor (e.g. press the sensor which opens the mouth), translate verbally the answer for the experimenter						P		P	P	
	Symbolic & pretend play Complex play with an accessory										
	Symbolic & pretend play Complex nap with Aibo										
	Symbolic & extrapolation play : "RobotCat" - Speak about the idea of a robotic cat (possibly imagine how one would play with it)										
	Causal composition of plays: Bite/Save & Give Food/Drink							P			r
	Causal composition of plays: Kiss & Bite/Save										
	Pretend play & causal reaction & social rapports: Ask verbally Aibo to act a situation, use of sensors									P	
L 6	Pretend play & focus on Aibo's mental states: Mimic Aibo's cry, and explain Aibo is never crying but pretending to cry										
	Pretend play & social rapports: Look after Aibo and set up rules										
	Pretend & symbolic & chronological play & social rapports: Search and rescue										
	Pretend & symbolic play & social rapport & cause-reaction play & chronological play: competition (drink fast) between the child or the experimenter and Aibo ; the non-competitor activates Aibo's sensor							P			

Figure 17. Child E. Play Grid. See Fig. 11 for a detailed caption

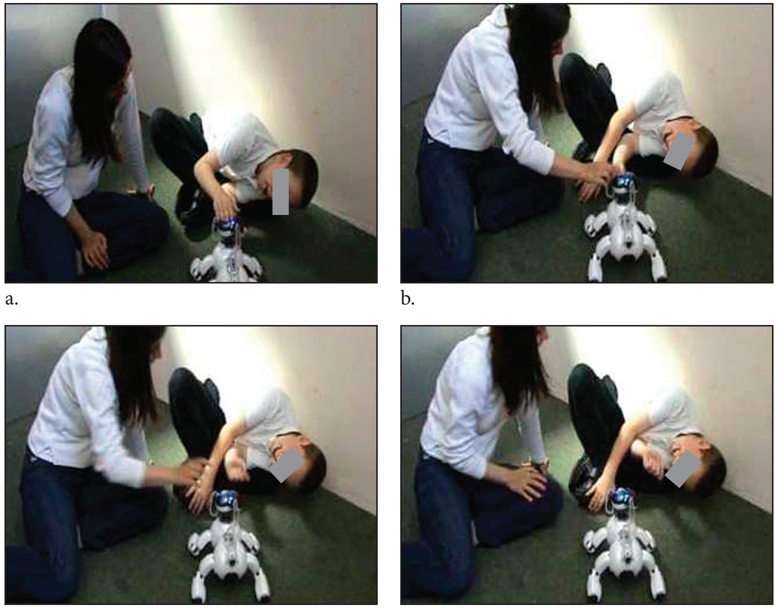

a. b. c. d.

Figure 18. Child E playing the game 'Bite/Save' with the experimenter. Chronological order of the photos: from left to right and top to bottom. a) the child activates the head sensor of the robot which makes the robot open the mouth and enables the robot to 'bite' his finger. b) the experimenter brings her hand close to the head of the robot in order to activate the head sensor. c) the experimenter activates the robot's head sensor to make Aibo open the mouth in order to 'save' the child's finger; when the mouth opens, the child pulls of his finger (c and d)

or the experimenter (through the use of the sensors) and (ii) the person remaining (child or experimenter) saved the latter by freeing her/his finger: the freeing was done either by activating the sensor ("Complex Bite/Save") or by directly taking the finger out of the mouth of the robot ("Simple Bite/Save").

Furthermore, in Session 7, the child proactively combined 2 games, "Give food/drink" and "Bite/save" and said: "He (the robot) is saying: give me a drink or I bite your fingers".

Another interesting play situation the child proactively experienced in Session 7 consisted of a competition between the robot and himself: both of them had to drink as fast as possible their invisible drink; the robot could only drink with the help of the experimenter (the experimenter was asked to activate the sensor linked to the opening

of the mouth as fast as possible). At the end of the competition, Child E decided that the robot had won. Thus, in this play situation Child E experimented with:

a. dealing with rules of competition,
b. handling the temporal aspects of the game and the various chronological phases,
c. taking on the role of the participant (as a competitor) and the one of the organizer who announces the winner,
d. playing with abstract entities (invisible drink),
e. playing socially.

Concerning the reasoning dimension, it should be first noted that the child decided to rename the robot after the first session and call him "Tornado". Moreover, in the first sessions, most of his questions addressed the issue of the robot's technical capabilities and how to control the robot. In Session 2, for instance, the child said: "How is he doing that?" and "What's being on the head to make him walk?" (because when he touched the head and activated the head sensor, the robot walked). And later in the same session, while looking at the laptop he said "this must be the controller". Furthermore, in Session 3, the child said: "I found how he might open his mouth"; the experimenter asked "is he moving the mouth?" and the child answered: "yes, when I stroke on the head, you see". This example illustrates that the child actively developed technical and causal reasoning about behaviours and capabilities of the robot. This questioning can be related to the category "Essence" and shows that the child considered primarily Aibo (Tornado) as a proper robot. It should be noted here that the child invented the concept of "invisible drink" as well as the way of calling it (very logically): "invisible robot drink". This illustrates the ability of the child to make links with a real dog's life while adapting it correctly to the characteristics of robots.

The category "Mental state" was addressed during later sessions (from Session 5 onwards). In Session 5 the child actually said "he is wagging the tail"; the experimenter answered: "yes, that shows he is happy"; and the child replied "He likes me" and he stroked the robot. The experimenter reinforced the positive feeling: "yes, he likes you". That first step was expanded into the game "speak directly to Aibo about Aibo's feeling". In Session 6 and onwards, the child addressed proactively the question of emotions but he tended to deal with a restricted repertoire of emotions only, such as "being scared" or "being terrified" (e.g. Session 7 the child said: "You're scared Tornado, in fact you're terrified").

Child E dealt with "Moral standing" in Session 5 when he accidentally kicked the robot and, in return, apologized to it directly ("Sorry Tornado") and comforted it by stroking him.

Finally, Child E addressed indirectly the question of "Social rapport" through play. For instance, in Session 10, he conferred a specific role to the robot for the competition; the robot thus became his adversary, but on a very kind level, since the child decided at the end of the game that the robot had won the competition. Another example took place in Session 8 where the child asked directly questions to the robot (e.g. "Do you want to drink something Tornado?"). Then, he made the robot bark as an answer and the child "translated" the answer verbally for the experimenter: "He said yes". In this case, the child proactively played the social role of an intermediary position between the experimenter and the robot.

The dimension of affect (Fig. 19) appeared from Session 5 onwards where the child proactively said "he (the robot) likes me". And the experimenter replied "Yes he likes you. You like him?" The child then answered "Yes". Then later, in Session 8, the child said "he (the robot) is very happy". The experimenter agreed with him and then Child E added "Tornado likes me" and the experimenter reinforced the positive feeling: "Yes he likes you". In Session 9, Child E commented on the robot, qualifying him as "friendly": "Tornado is very friendly, isn't it?" and the experimenter agreed verbally.

Child F. Child F was away for Session 5. Thus he took part in 9 sessions. Note that on his explicit demand, Session 7 and Session 8 were not recorded (the experimenter had permission from the parents to videotape the child but she decided to value the child's request); thus information from sessions 7 and 8 is missing in the corresponding columns in the Play Grid. The Play Grid (Fig. 20) shows that

Session	**Events objectively related to Affect** (ordered chronologically with respect to first appearance, event only mentioned once per session)
S1	
S2	
S3	
S4	
S5	· **[1i]** "Yes that shows he (the robot) is happy" (Experimenter); **"He likes me"** (Child E); "Yes he likes you" (Experimenter); · **[2i]** **"You like him** (the robot)**?"** (Experimenter); **"Yes"** (Child E)
S6	
S7	
S8	· **[1i]** "He (the robot) is very happy" (Child E) while making the robot bark; "Yes he is" (Experimenter)**, "Tornado likes me"** (Child E); "Yes he likes you" (Experimenter)
S9	· **[1ii]** **"Tornado is very friendly, isn't it?"** (Child E); "yes, he is"(Experimenter)
S10	

Figure 19. Child E. Events related to Affect. See caption of Fig. 12 for details

		1	2	3	4	5	6	7	8	9	10
L 1	*Solitary* Exploration										
	"Imitation" of robot's bark	P	P	P			P			P	
	Solitary mirror play – look at oneself in the robot's reflecting face										
L 2	*"Pre-social" or basic-social* exploration – stroke Aibo immediately after the experimenter (possibly basic imitation of the gesture)										
L 3	*Social* exploration (social play)	P	P	P	P		P			P	P
	Simple Bite/Save or Give/Food - no use of the sensors									P	P
	Position or locomotion game – with verbal qualification of the game	P			P		B			B	P
	Cooperative technical task: change the battery, or turn on/off Aibo	r		P	B		r			P	B
	Verbal order towards Aibo: e.g. "sit", "walk", "wake up"	P	P	P	P					B	P
	Basic pretend & social play – imitate Aibo's snoring & verbal comment										
	Basic play on affective gestures – give/receive a kiss and/or a lip to/from Aibo									P	P
	Repeat after me - ask the experimenter to repeat verbal expressions										P
	Look at Aibo through the camera (Possibly stroke Aibo & look at its reaction through the camera)			P	P		P			P	P
	Speak French with Aibo - e.g. "Hello" or "Bye-Bye" in French				r		B				r
	Show Aibo to other children (social play) Express verbally the willing/intention to show Aibo to the other children	P	P								
	Simple play with accessory (symbolic play)			P	P						
	Social Mirror play (social play) - look at oneself (and possibly at the experimenter) in the robot's reflecting face & express verbal comments, e.g. "Look at my arm!"										
	Social Hug – hug Aibo & ask the experimenter or the second researcher to hug Aibo										
L 4	*Complex* Give Food/Drink (cause-reaction play & symbolic play & social play) - use of sensors										
	Complex Bite/Save (cause-reaction play & pretend play & cooperative play) - use of sensors										
	Complex turn off Aibo to sleep (symbolic play)						P				P
	Speak directly to Aibo about Aibo's feeling (symbolic play)		P								
	Cause-reaction play & mental states: Ask a question to Aibo (e.g. identity, feeling), answer with a sensor		B	P	r		B				
	Cause-reaction play, Aim at a physical reaction of the robot, show it with a sensor		P	B	B		r			P	P
	Cause-reaction play & basic pretend play, "caught on the act"										
	Telling a story										
L 5	Cause-reaction play and explicit Social rapport: Ask a question to Aibo, answer with a sensor (e.g. press the sensor which opens the mouth), translate verbally the answer for the experimenter										
	Symbolic & pretend play Complex play with an accessory			P	P		P				
	Symbolic & pretend play Complex nap with Aibo				P						
	Symbolic & extrapolation play : "RobotCat" - Speak about the idea of a robotic cat (possibly imagine how one would play with it)									P	P
	Causal composition of plays: Bite/Save & Give Food/Drink										
	Causal composition of plays: Kiss & Bite/Save										P
	Pretend play & causal reaction & social rapports: Ask verbally Aibo to act a situation, use of sensors										
L 6	Pretend play & focus on Aibo's mental states: Mimic Aibo's cry, and explain Aibo is never crying but pretending to cry										P
	Pretend play & social rapports: Look after Aibo and set up rules										
	Pretend & symbolic & chronological play & social rapports: Search and rescue										
	Pretend & symbolic play & social rapport & cause-reaction play & chronological play: competition (drink fast) between the child or the experimenter and Aibo ; the non-competitor activates Aibo's sensor										

Figure 20. Child F. Play Grid. See Fig. 11 for a detailed caption. Note that Child F was away for Session 5 and, on his request, was not filmed during Sessions 7 and 8

Child F engaged in social play almost all the time. He used verbal language a lot and progressively experienced some more complex levels of play, notably pretend play with respect to "play with accessory". The first situations of "play with accessory" happened in Session 3. In this session, the child borrowed the mouse of the laptop and put it on the ground in front of Aibo at approximately 30 cm distance and asked the robot to touch the mouse with the paw. Then he activated the right sensor to make Aibo walk forward and approach the mouse. The child carried the robot for the 5 remaining centimetres separating the robot's paw from the mouse and finally the robot touched the mouse with his paw. Later, in Session 4, the child experienced further situations of "play with accessory" in two successive steps. As a first step, he proactively played very simply with an accessory. For instance, Child F used the face of a character drawn on a piece of cardboard that he held in front of his face and told Aibo: "Stay here Aivo, I've got something to show you". Note that the child slightly changed the pronunciation of the name of the robot and referred to Aibo as 'Aivo'. As a second step, later in the same session, the child proactively played a more complex accessory game with the robot, the "ghost dog". That play situation consisted in putting a cloth on top of Aibo and pretending Aibo was a ghost dog (Child F told Aibo: "You can be a ghost dog Aivo"); vocally, the child used classical onomatopeia mimicking a ghost's "voice and presence". Moreover, in Session 6, the child decided to make the robot wear clothes and this game was expanded by:

a. a series of questions on inferring states of the robot with respect to like/dislike,
b. a direct communication with the robot to explain to it what he was wearing (Child F told Aibo: "Look at you Aivo! You've got some paper on to be black");
c. a version of the game "aim at a physical reaction of the robot, show it with a sensor" (the experimenter asked "How do you make him walk with all these clothes?", the child replied "Walk?", and the child made the robot walk).

In addition to the accessory games, the child experimented with pretend play with the robot in a social context, e.g. pretending to have a nap with the robot (in Session 4) in a detailed (and complex) way resulting in

1. using a cloth as a blanket to cover both of them,
2. deciding on the duration of sleep and asking for the clock to be watched to respect the time predefined for the nap,
3. pretending to snore,
4. both of them waking up again.

Besides, another way of tackling pretend play as well as the robot's mental states happened in Session 10 when the child imitated Aibo's crying, and then argued that Aibo was not crying but pretending to cry. And this notion of pretending to cry for the robot was reused many times during the last Session (e.g. Child F said: "No, he's not crying, he is only pretending to cry.").

The reasoning dimension is an important component of the profile of Child F. Child F principally addressed three of the four components: "Essence", "Mental States" and "Social Rapport", and, to a lesser degree, "Moral Statement".

Concerning "Essence", the child really tackled the question of artefact or biological features, processes and categories. In relation to category, he often asked about the robot dogs' boundaries, e.g. in Session 2: "Have you seen dogs that are not robot dogs, yes or no?" he asked the experimenter, and later in the same session: "He has short teeth, he doesn't bite. Robot dogs don't bite, do some do?"

The part on "Mental States" component is very rich since the child addressed all the aspects defined in the coding manual of Kahn et al. (2003) except probably the "autonomy" one. Actually, he attributed "intentions" to the robot in Sessions 1 and 2. He explicitly considered the robot's "emotional states" in sessions 2, 4, 6 and 10. He also both tackled "emotional states" of the robot and his "personality" when he asked the robot questions about its likes/dislikes (e.g. Session 4: "Do you like toys Aivo, yes or no?"). Furthermore, he pretended the robot had some "cognitive abilities" and developed play upon it: in Session 4, for instance, he disguised himself with an accessory in order to "show" Aibo and thus presupposed -for the game- that Aibo could see. Later, in Session 6, again the child presupposed for the game that the robot could see and told it: "Look at you Aivo. You've got some paper on to be black". The last aspect of "mental states" is the notion of "development" of the robot. Child F asked about it throughout the sessions. More than the notion of development, the child seems to have been willing to build a biography for the robot (i.e. the past of the robot) and therefore asked questions to the experimenter such as: (a) in Session 1: "Where was this robot dog from?"; (b) in Session 2: "Where was he born?" and "Has he travelled in a car?"; (c) in Session3: "Where did you get him from?", "Where does he live?", "How old is he?", etc.

Concerning the part on "Social rapport", the child really investigated the social links between the robot and the experimenter, who was considered by the child as being the "mum" of the robot (Child F told the experimenter "it's your dog son", meaning that Aibo is the experimenter's dog, and that the experimenter, in a way, is considered as being Aibo's 'mum'). He also investigated the social links between the robot and himself, through situations of pretend play but also verbally. In Session 2 for instance, the child presupposed that there was a social rapport between the robot and himself since he told the robot: "When it is lunch time Aivo I got to go. And don't cry Aivo". Later, in Session 6, the child stated that the robot was his cousin: "Aivo is my cousin". And when the experimenter asked: "Aivo, do you like playing with F?[16] Can you tell me? Can you ask for his answer F?" then the child told Aibo: "Aivo do you like me? You're my cousin. I'm your cousin, Aivo". The child also investigated beyond social rapport involving Aibo and, for instance, asked the experimenter a few questions about her family: (a) in Session 4, the child asked about the experimenter's French accent:[17] "What accent do you speak?", which was further investigated in Session 6: "Why do you speak French?" and "Why were you born in France?"; (b) in Session 6, he asked her about her family: "What are your parents' names?"; he investigated further questions on the experimenter's family in Session 10.

On the "Affect" level (Fig. 21), the child expressed himself a lot, both by gestures (e.g. giving a kiss to Aibo after saying "Goodbye Aivo, have a good sleep" in Session 6) and verbal expressions (e.g. in Session 4 when he dressed up Aibo: "Put this on, Aivo, my dog, my friend, Aivo"). It is perhaps worthy of note here that it might be the case that some gestures relating to affect from a non-autistic perception (e.g. giving a kiss), do not have the same interpretation for a child with autism: for a child with autism, giving a kiss might, for instance, just be an imitated response. Concerning Child F, it might be the case that the child reproduced the gesture "giving a kiss" from a situation he had encountered or witnessed before; nonetheless it should be mentioned that his gesture was made proactively, with no previous reference from the experimenter to such a gesture.

Session	**Events objectively related to Affect** (ordered chronologically with respect to first appearance, event only mentioned once per session)
S1	· **[1ii]** **"Ooh he is a nice dog"** (Child F) and he strokes the robot
S2	
S3	
S4	· **[1ii]** Child F brings a towel to put on the robot : **"Put this on Aivo, my dog, my friend, Aivo"** (Child F)
S5	
S6	· **[1i]** **"Aibo, do you like me?** You're my cousin. I'm your cousin Aivo " (Child F) · **[1iv]** Child F gives a **kiss** to the robot on the muzzle after saying **"OK, Goodbye Aivo, have a good sleep"**
S7	
S8	
S9	
S10	· **[1iv]** Child F has covered Aibo with a coat; he gives the robot **a kiss** on the forehead and says **"Goodnight Aivo"**

Figure 21. Child F. Events related to Affect. See caption of Fig. 12 for details

6. Discussion

Results from these experiments show that the children progressed differently, and that their profiles according to the three (intertwined) dimensions *Play – Reasoning – Affect* are unique. This highlights how the experimental approach presented in this study allows many trajectories for progressing and, more specifically, how it can meet the child's specific needs and abilities.

Furthermore, concerning the dimension of play, and, more precisely, concerning the children's progression with respect to solitary vs. social play, three groups can be highlighted. The first one, group 1, consists of children who mostly played solitarily and possibly encountered rudimentary situations of imitation, but no further components of social play. This group includes Child A who encountered imitation in Session 10 and Child B. Note that both of them find it very hard to communicate verbally. For the children whose current play with the robot is mainly dyadic, it is particularly relevant to enable the robot to adapt automatically to their play styles in real time so that they can benefit from this dyadic play and progressively reach well balanced and potentially higher levels of play. The second group, group 2, consists of Child D who communicated mainly non-verbally yet progressively experienced situations of verbal communication and showed pre-social or basic social play during the last sessions. The third group, group 3, consists of Child C, E and F. Those children proactively played socially (i.e. in a triad including both the robot and the experimenter).

For those three groups, results shows that a) Child B (group 1) experienced progressively longer uninterrupted periods of play and engaged in basic imitation during the last sessions; (b) children from group 3 tended to experience higher levels of play gradually over the sessions and constructed more and more reasoning about the robot (and sometimes engaged in specific reasoning about real life situations as well). At a more basic stage, Child D (group 2) also experienced higher levels of play progressively. He started to reason about technical aspects of the robot as well, e.g. 'turning on/off' the robot and changing the battery. In the last sessions different elements suggested that he may also have experienced some reasoning about social rapport. Besides, the children's proactivity was encouraged, enabling them to take initiative and express intentions (cf. the proportion of proactive activities vs. reactive activities in the Play Grids).

These results are in agreement with Josefi et al.'s findings (cf. Section 'Related Work') who have shown that non-directive play therapy encouraged the child's initiative-taking (Josefi & Ryan, 2004). Further to this, Josefi et al.'s study has shown that non-directive play therapy may encourage symbolic play, which is an important finding of our approach too: In our study, children from group 3 progressively experienced situations of symbolic or pretend play. Note that, as already explained, the study presented here took place in a therapeutic context but the experimenter was not behaving exactly like a therapist.[18] Besides, we identified several advantages in introducing a robotic pet in the experimental setup:

a. the use of a robot allows us to simplify the interaction and to initially create a relatively predictable environment for play, thus facilitating the child's understanding of the interaction (e.g. by initially giving the robot a simple predictable behaviour) (Dautenhahn & Werry, 2004). Progressively the complexity of the interaction can be increased.
b. children tend to express interest in the robot, and occasionally affect towards Aibo, as our findings show;
c. here, one of the findings is that, in these experiments, with this new approach, through play with the robotic pet, children tend to develop reasoning, and make comparisons to real dogs' lives. Note that based on our findings we cannot claim that the children's reasoning genuinely developed as a direct result of our study – we observed, however, cases where reasoning skills were *expressed* increasingly during successive sessions. Thus, the robotic pet can be considered as a good medium for developing and/or expressing reasoning on mental states and social rapport upon, and for learning about basic causal reactions.

In the context of robot-assisted play, we have shown in Section 'Related Work' that research has, until now, mainly addressed task-oriented activities, such as chasing games with Labo-1 (Werry & Dautenhahn, 1999) or imitation with Robota (Robins et al., 2004). Nadel et al. have shown that imitation skills have a significant impact on the acquisition of social skills for children with autism (Nadel et al., 1999). However, focusing on imitation tasks only may not be sufficient when the child reaches some higher levels of play (cf. children from group 3 in the experiments presented in this study); Howlin and Rutter underlined the necessity of incorporating developmental aspects (Howlin & Rutter, 1987).

The study presented in this paper goes beyond these previous experiments, since it provides the child with a relatively highly unconstrained environment of play: due to the mobile and autonomous nature of the robotic pet, the child can engage in a larger repertoire of play situations (note that Robota was remotely controlled and fixed in place while Labo-1, while operating autonomously, had no tactile sensors) and notably experience causal reaction play and symbolic play. Imitation is used to bootstrap and initiate more complex situations of interaction or

to help the child re-engage in the interaction. Besides, in this approach the experimenter is both a "passive participant" and, under precise conditions, becomes an active participant, which expands and formalizes his/her role compared with Robins et al.'s study (Robins & Dautenhahn, 2006).

Moreover, in this study, we have adopted a qualitative approach for the analysis of each dimension, Play, Reasoning and Affect. We were actually interested in the emergence and in the specificities of the play styles, questions or statements related to reasoning and events that could be objectively related to affect, rather than in the occurrences or the duration of each of them. In particular, two similar games might actually happen to be different in the way the child experiences them, such as for example, the fluency, the rhythm, the coherence etc. Consequently, unlike a quantitative analysis which often relies on micro-behaviour analyses[19] (e.g. Dautenhahn & Werry (2002); Tardif et al. (1995)), this qualitative analysis here focused on a bigger scale, i.e. an intermediary scale.[20] This intermediary scale enabled us to consider events constituting a game as connected events and, in particular, to describe the structure of a specific play situation in possibly different (chronological) phases or identify in this play situation, the presence of social play, the proportion of symbolic or pretend play, and the use of causality.

This study is explorative in nature, and more research should be done to investigate more systematically the contribution of such an approach in the field of robot-mediated therapy for children with autism.

7. Future work

Looking back at the results, the existence of group 1 shows that some children remained playing mainly dyadically with the robot. The only situations of social play those children experienced were basic imitation. For those children, it is particularly crucial to develop basic play skills through this dyadic interaction first, in order to help them reach higher levels of play and ideally, experience later triadic situations of play with the experimenter and the robot.

As part of future work, the question should therefore be investigated as to how to further facilitate children's play with the robot, for the children who remain at the level of solitary play; in this case, the robot should be able to adapt appropriately to the child's needs and abilities and encourage the child to progress towards more complex play styles autonomously. This issue has been addressed in François et al. (2007, 2008) where the robot adapts its behaviour in real time and autonomously to specific play styles of the child in order to guide him/her towards more balanced interaction styles. Such an 'adaptive' robot might help the child e.g. experiment with simple cause-reaction play situations. In François et al. (2009) we implemented and evaluated such an adaptive robot that rewards well-balanced interaction styles (e.g. not too strong, not too frequent) in a study conducted with seven children with autism. A statistical analysis of the results showed the positive impact of such an adaptive robot on the children's play styles and on their engagement in the interaction with the robot. Such initial findings are promising and need to be extended in future larger-scale studies.

Ideally, at some point, the child would naturally move towards group 2 and be able to engage in simple situations of social play (with both the experimenter and the robot).

Another avenue for future research within the proposed approach is to include "theory of mind" (ToM) more explicitly in the experimental design. ToM was not specifically considered in the present work, but we observed children commenting on the robot's intentions and 'feelings', which may provide a starting point for more detailed ToM investigations. Children with autism's difficulties with "mindreading" have been reported widely (e.g. Baron-Cohen et al. (1985); Hobson (1993); Baron-Cohen (1997)) and its relevance to the employment of interactive robots in autism therapy has been discussed in Dautenhahn & Werry (2004). Thus, in the context of the approach presented in this article, future work could specifically include further aspects of ToM, e.g. concerning the children's abilities to read the experimenter's intentions, goals and beliefs, or to take his/her perspective during play. Also, future work in this area would benefit from an assessment of whether the skills – that the children developed and/or expressed during the play sessions according to our approach – will also generalize to other situations, e.g. involving other children or adults (instead of the experimenter), or involving other play and/or social interaction situations within and outside the school. Furthermore, in order to distinguish whether the skills expressed by the children in our sessions genuinely developed or whether our approach only helped them to better express them in successive sessions, an assessment and comparison of the children's skills prior and after the play sessions in different contexts are important. Such directions would benefit from a larger-scale research programme put together and carried out jointly by roboticists, autism researchers as well as therapists, teachers and possibly also involving the children's parents.

Generally, future work in this area could either encompass more parameters to test, e.g. include further specific aspects of ToM as discussed above, or it could concentrate in further depth on specific aspects such as the dimension of "Play" and e.g. investigate in great detail different levels and aspects of play.

8. Conclusion

This paper presents a more approach in the context of robot-mediated therapy with children with autism. This approach draws inspiration from non-directive play therapy, notably encouraging the child's proactivity and initiative-taking. Here, the experimenter participates in the play sessions and the child is the main leader for play. Beyond inspiration from non-directive play therapy, the approach introduces a regulation process: the experimenter can regulate the interaction under specific conditions; in brief:

a. to prevent or discourage repetitive behaviours,
b. to help the child engage in play,
c. to give a better pace to the game if it has already been experienced by the child,
d. to bootstrap a higher level of play,
e. to ask questions related to reasoning or affect.

A long-term study was carried out with six children which highlighted the capability of the method to adapt to the child's specific needs and abilities through a unique trajectory of progression with respect to the three dimensions, Play-Reasoning-Affect. In particular, each child made progress in at least one of the three dimensions progressively over the sessions. Moreover, in terms of play, and, more precisely, solitary vs. social play, children could be categorized into three groups. The children who managed to play socially experienced progressively higher levels of play and developed progressively more reasoning related to the robot; they also tended to express some interest towards the robot, including on occasions interest involving positive affect. This preliminary long-term study has therefore shown promising results for this new approach in robot-assisted play. It is a first study that potentially may be developed towards a new method in autism therapy.

Acknowledgments

The first author is supported by a research scholarship of the University of Hertfordshire. The work described in this paper was partially conducted within the EU Integrated Project RobotCub (Robotic Open-architecture Technology for Cognition, Understanding and Behaviours) and was partially funded by the European Commission through the E5 Unit (Cognition) of FP6-IST under contract FP6–004370.

The authors would like to thank the teachers at the school, the children who took part in the experiments, and their parents. The authors would like to thank Megan Davis for assisting in some of the experiments.

Notes

1. http://www.nas.org.uk/

2. It should be further noted that children with autism often tend to perceive objects in their parts and not as a whole, which is integral to the weak central coherence theory (Fritz, 1989). This frequent inability may also influence the way the child plays.

3. Onomatopoeia refers to using words that imitate the sound(s) associated with objects or actions, e.g. "buzz".

4. We focus on Rogerian theory in this article because of its strong ties with non-directive play therapy which is a key source of motivation for us. Our specific approach to robot-assistive play as outlined in this article is however not restricted solely to his theory. Other theoretical approaches such as those proposed by Jerome Bruner (Bruner, 1986, 1990) and Lev S. Vygotsky (Vygotsky, 1978) – both are indeed fundamental to other work in our research group related to development, narrative and learning – could be used for an extended theoretical discussion of this work that would however go beyond the scope of this publication.

5. "(i) emotional security and relaxation, (ii) an enhanced and attentive adult environment in which playing together is emphasized, and (iii) the acceptance by therapists of children's ability to instigate therapeutic change for themselves under favourable conditions". (Josefi & Ryan, 2004:545).

6. Note, the symbolizing capacities have similarities with, and may overlap with, capacities to learn language during normal development; conversely, it is very likely that learning a language requires some symbolizing capacities and processes.

7. The seal robot Paro was introduced in the Bobath protocol (http://www.bobath.org.uk/) in the context of a child with severe cognitive and physical delays. The Bobath protocol is a method used for the rehabilitation of physical functional skills (Knox & Evans, 2002). Results showed that the introduction of Paro may have strengthened, for this particular child, the efficiency of the Bobath protocol.

8. Those imitations concerned the position or movement of arms and legs.

9. In this study, focused shared attention refers to the child's eye gaze directed towards the mediator (alternatively a human or a robot). It does not include joint visual attention, i.e. looking at an object that the mediator is pointing at.

10. The experimenter was the first author of this paper.

11. Note, the ultimate goal of this approach is to prevent the child from exhibiting repetitive behaviour in the first place.

12. Different classifications of play coexist in play literature. Piaget's classification identifies four categories: practice play, symbolic play, games with rules and constructions (Piaget, 1945). Another

taxonomy, given by Boucher (1999) emphasizes the importance of social play which is one category of the classification. Here, we take a slightly different perspective since we *describe (analyse)* each situation of play according to four criteria, which are, in this context of robot-assisted play for children with autism, of particular relevance to measuring progress in the expression of skills in social interaction, communication, reasoning related to the robot and imagination. A situation of play is analysed according to the four criteria. These criteria are not exclusive to each other. On the contrary, a situation of play should ideally contain several of these criteria.

13. Shanti is the name of the stuffed dog that was used in Kahn et al. (2003)'s study as a basis for comparison.

14. The recoded segments contained only high involvement of the children in interaction. High involvement is characterised by the fact that (i) children do not stop interacting for a period longer than a few seconds, and (ii) children experience many situations of play, reasoning or affect related to the robot. Therefore, the density of events to identify and code is very high in the recoded segments which makes the evaluation highly meticulous.

15. For instance, the experimenter did not ask the question: "Does the robot hate you?", to which the child might have said "yes" as well.

16. Child F is designated by F in the dialogue.

17. Child F mastered some French vocabulary.

18. The experimenter did not have any formal training as a therapist.

19. Micro-behaviour analysis is the analysis of videos based on the coding of low level behaviours such as eye gaze, eye contact, touch, etc.

20. To make a parallel with the notion of micro-analysis used in (Tardif et al., 1995), one could qualify our approach here as a mesoscopic approach or a meso-analysis. The prefix 'meso' comes from the Greek word 'mesos', meaning middle. "Mesoscopic" is an intermediary scale between "microscopic" and "macroscopic". Those terms are commonly used in physics and chemistry, and can be transposed metaphorically to our context. Applied to our context here, a mesoscopic approach means that we look at the events constituting an uninterrupted game as connected events, and as a whole.

References

Association, American Psychiatric. (1994). *Diagnostic and Statistical Manual of Mental Disorders DSM-IV, fourth edition*. American Psychiatric Publishing.
Aurora (2009). http:/www.aurora-project.com. Last accessed 31/10/2009.
Axline, V. (1946). *Dibs: In search for self*. New York: Ballantine Books.
Axline, V. (1947). *Play therapy*. New York: Ballantine Books.
Baillie, J.-C. (2005). URBI: Towards a universal robotic low-level programming language. In *Proc. IEEE/RSJ Int. Conf. on Intelligent Robots and Systems* (pp. 820–825).
Baron-Cohen, S. (1997). *Mindblindness, an essay on autism and theory of mind.* The MIT Press, new edition.

Baron-Cohen, S., Leslie, A.M., & Frith, U. (1985). Does the autistic child have a 'theory of mind'? *Cognition*, 21(1), 37–46.
Boucher, J. (1999). Editorial: Interventions with children with autism – methods based on play. *Child Language teaching and therapy*, 15, 1–5.
Bruner, J. (1986). *Actual minds, possible worlds*. Cambridge, Mass.: Harvard University Press.
Bruner, J. (1990). *Acts of meaning*. Cambridge, Mass: Harvard University Press.
Chaillé, C., & Silvern, S. (1996). *Understanding through play*. The British Library – The World's Knowledge".
Dautenhahn, K. (2007). Socially intelligent robots: Dimensions of human–robot interaction. *Philosophical Transactions of the Royal Society B,* 362(1480), 679–704.
Dautenhahn, K., & Werry, I. (2000). Issues of robot–human interaction dynamics in the rehabilitation of children with autism. In *Proc. From Animals to Animats, 6th Int. Conference on the Simulation on Adaptive Behavior (SAB2000)* (pp. 11–15).
Dautenhahn, K., & Werry, I. (2002). A quantitative technique for analysing robot–human interactions. In *Proc. IEEE/RSJ Int. Conference on Intelligent Robots and Systems* (pp. 1132–1138).
Dautenhahn, K., & Werry, I. (2004). Towards interactive robots in autism therapy: Background, Motivation and Challenges. *Pragmatics and Cognition*, 12(1), 1–35.
Duquette, A., Michaud, F., & Mercier, H. (2007). Exploring the use of a mobile robot as an imitation agent with children with low-functioning autism. *Autonomous Robots*, 24(2), 147–157.
François, D., Dautenhahn, K., & Polani, D. (2009). Using real-time recognition of human–robot interaction styles for creating adaptive robot behaviour in robot-assisted play. In *Proc. Second IEEE Symposium on Artificial Life.* (pp. 45–52).
François, D., Polani, D., & Dautenhahn, K. (2007). On-line behaviour classification and adaptation to human–robot interaction styles. In *Proc. 2nd ACM/IEEE International Conference on Human–Robot Interaction (HRI07)* (p. 295–302).
François, D., Polani, D., & Dautenhahn, K. (2008). Towards socially adaptive robots: A novel method for real time recognition of human–robot interaction styles. In *Proc. IEEE-RAS International Conference on Humanoid Robots (Humanoids08)* (pp. 353–359).
Fritz, U. (1989). *Autism: Explaining the enigma*. Oxford, Blackwell.
Hobson, R.P. (1993). *Autism and the development of mind*. Lawrence Erlbaum Associates Ltd.
Howlin, P., & Rutter, M. (1987). *Treatment of autistic children*. John Willey and Sons.
Josefi, O., & Ryan, V. (2004). Non-directive play therapy for young children with autism: A case study. *Clinical child Psychology and Psychiatry*, 9(4), 533–551.
Kahn, P.H., Friedman, B., Freier, N.G., & Severson, R. (2003). *Coding manual for children's interactions with AIBO, the robotic dog – the preschool study* (UW CSE Technical Report 03604603). Department of Computer Science and Engineering, University of Washington.
Kahn, P.H., Friedman, B., Pérez-Granados, D.R., & Freier, N. G. (2006). Robotic pets in the lives of preschool children. *Interaction Studies*, 7(3), 405–436.
Knox, V., & Evans, A. (2002). Evaluation of the functional effects of a course of Bobath therapy in children with cerebral palsy: a preliminary study. *Developmental Medicine and Child Neurology*, 44(7).
Kozima, H., Nakagawa, C., & Yasuda, Y. (2005). Interactive robots for communication-care: A case study in autism therapy. In *Proc. 14th IEEE Int. Workshop on Robot and Human Interactive Communication (RO-MAN)* (pp. 341–346).

Loureiro, R., Amirabdollahian, F., Topping, M., Driessen, B., & Harwin, W. (2003). Upper limb robot mediated stroke therapy – GENTLE/s approach. *Autonomous Robots*, 15(1), 35–51.
Marti, P., Fano, F., Palma, V., Pollini, A., Rullo, A., & Shibata, T. (2005). My Gym Robot. In *Proc. AISB'05 Symposium on Robot Companion. Hard Problem and Open Challenges in Human–Robot Interaction* (pp. 64–73).
Nadel, J., Guerini, C., Peze, A., & Rivert, C. (1999). The evolving nature of imitation as a format for communication. In J. Nadel & G. Butterworth (Eds.), *Imitation in infancy* (p. 209–234). Cambridge University Press.
Piaget, J. (1945). *La formation du symbole chez l'enfant: Imitation, jeu et rêve, image et représentation*. Delachaux et Niestle.
Powell, S. (2000). *Helping children with autism to learn*. David Fulton Publish.
Robins, B., & Dautenhahn, K. (2006). The role of the experimenter in HRI research – a case study evaluation of children with autism interacting with a robotic toy. In *Proc. 15th IEEE Int. Workshop on Robot and Human Interactive Communication* (RO-MAN). (pp. 646–651).
Robins, B., Dautenhahn, K., Boekhorst, R. te, & Billard, A. (2005). Robotic assistants in therapy and education of children with autism: Can a small humanoid robot help encourage social interaction skills? *Universal Access in the Information Society (UAIS).* (pp. 105–120).
Robins, B., Dautenhahn, K., & Dubowski, J. (2005). Robots as isolators or mediators for children with autism? A cautionary tale. In Proc. AISB'05 *Symposium on Robot Companion. Hard problem and Open Challenges in Human–Robot Interaction* (pp. 82–88).
Robins, B., Dickerson, P., Stribling, P., & Dautenhahn, K. (2004). Robot-mediated joint attention in children with autism: A case study in robot–human interaction. *Interaction Studies*, 5(2), 161–198.
Rogers, C. (1976). Client-centred therapy. Constable.
Ryan, V. (1999). Developmental delay, symbolic play and non-directive play therapy. *Clinical child Psychology and Psychiatry*, 4(2), 167–185.
Ryan, V., & Needham, C. (2001). Non-directive play therapy with children experiencing psychic trauma. *Clinical child Psychology and Psychiatry*, 6(3), 437–453.
Ryan, V., & Wilson, K. (1996). *Case-studies in non-directive play therapy*. Bailli`ere Tindall.
Schatzman, L., & Strauss, A. (1973). *Strategy for recording*. Schatzmann, L., Strauss, A.L. Eds. Field Research. Strategies for a natural sociology. Prentice Hall.
Shibata, T., Wada, K., Saito, T., & Tanie, K. (2005). Human interactive robot for psychological enrichment and therapy. In Proc. AISB'05 *Symposium on Robot Companion Hard Problem and Open Challenges in Human–Robot Interaction* (pp. 98–109).
Stanton, C.M., Kahn, P.H., Severson, R.L., Ruckert, J.H., & Gill, B.T. (2008). Robotics animals might aid in the social development of children with autism. In *Proceedings of the 3rd International Conference on Human–Robot Interaction* (pp. 271–278).
Tanaka, F., Fortenberry, B., Aisaka, K., & Movellan, J.R. (2005). Developing dance interaction between QRIO and toddlers in a classroom environment: Plans for the first steps. In *Proc. 14th IEEE Int. Workshop on Robot and Human Interactive Communication* (RO-MAN) (pp. 223–228).
Tanaka, F., Movellan, J.R., Fortenberry, B., & Aisaka, K. (2006). Daily HRI evaluation at a classroom environment: Reports from dance interaction experiments. In *Proc. 1st Annual Conf. on Human–Robot Interaction (HRI)* (pp. 3–9).
Tardif, C., Plumet, M.-H., Beaudichon, J., Waller, D., Bouvard, M., & Leboyer, M. (1995). Microanalysis of social interactions between autistic children and normal adults in semi-structured play situations. *International Journal of Behavioural Development,* 18(4), 727–747.

Vygotsky, L.S. (1978). *Mind in society*: Development of higher psychological processes. Cambridge, Mass: Harvard University Press.

Werry, I., & Dautenhahn, K. (1999). Applying mobile robot technology to the rehabilitation of autistic children. In *Proc. SIRS99, 7th Int. Symposium on Intelligent Robotics Systems* (pp. 265–272).

Werry, I., Dautenhahn, K., & Harwin, W. (2001). Evaluating the response of children with autism to a robot. In Proc. RESNA 2001, *Rehabilitation Engineering and Assistive Technology Society of North America.*

Werry, I., Dautenhahn, K., Ogden, B., & Harwin, W. (2001). Can social interaction skills be taught by a social agent? The role of a robotic mediator in autism therapy. In *Proc. 4th Int. Conference on Cognitive Technology: Instruments of Minds (ct2001)* (pp. 57–74).

Author's address

Dorothée François
Adaptive Systems Research Group
School of Computer Science
University of Hertfordshire
College Lane
Hatfield, Herts, AL10 9AB
United Kingdom

Email:dorotheefrancois@hotmail.com

Stuart Powell
School of Education
University of Hertfordshire
College Lane
Hatfield, Herts, AL10 9AB
United Kingdom

Email:s.d.powell@herts.ac.uk

Kerstin Dautenhahn
Adaptive Systems Research Group
School of Computer Science
University of Hertfordshire
College Lane
Hatfield, Herts, AL10 9AB
United Kingdom

Email:k.dautenhahn@herts.ac.uk

About the authors

Dorothée François has a Ph.D. degree in Computer Science from the University of Hertfordshire, UK. Her research interests include Human–Robot Interaction, Robot-mediated Therapy and Education and Socially Adaptive Robots. She is involved in the Aurora Project and her Ph.D. research concerned facilitating play between children with autism and an autonomous robot. She has conducted play sessions in schools with children with autism for more than a year on a weekly basis, whereby she evaluated new approaches she designed for Robot-Assisted Play. In parallel, her research has addressed the issue of automatic real-time recognition of human–robot interaction styles.

Prof. Stuart Powell is Professor of Educational Psychology. His original research interests were in psychology and education, and he has published over a number of years in the area of special

needs teaching with a particular focus on working with children with autism. More recently, he has published in the areas of special needs and higher education and advanced study skills.

Prof. Dr. Kerstin Dautenhahn received her Ph.D. degree from the Biological Cybernetics Department of the University of Bielefeld (Germany) in 1993. She is Professor of Artificial Intelligence in the School of Computer Science and coordinator of the Adaptive Systems Research Group at the University of Hertfordshire (UK). She has published more than 100 research articles on social robotics, robot learning, human–robot interaction and assistive technology. Prof. Dautenhahn has edited several books and frequently organizes international research workshops and conferences. She is involved in several European robotics projects and is Editor in chief of the journal *Interaction studies.*

[Accepted Extended Abstract for
Full Proceedings Paper]

As Time Goes by: Representing and Reasoning Timing in the Human-Robot Interaction Studies

Hatice Kose-Bagci, Frank Broz, Qiming Shen, Kerstin Dautenhahn, Chrystopher L.Nehaniv
Adaptive Systems Research Group, University of Hertfordshire, School of Computer Science, Hatfield, Herts, AL 10 9AB, U.K.
{h.kose-bagci;f.broz;q.shen;k.dautenhahn;c.l.nehaniv}@herts.ac.uk

Abstract

We summarise the experimental design issues related to timing in three human-robot interaction studies investigating imitation and drumming experiences with child-sized humanoid robots and human participants. Our aim[1] is not to have the humanoid robots just replicate the human's behaviors (e.g. waving or drumming), but to engage in a 'social manner', i.e. in a call and response turn-taking interaction. This work is part of a research project on developmental robotics with a particular emphasis on imitation and gesture communication.

1. Introduction

Timing plays a fundamental role in the regulation of human-robot interaction and communication. We present the experimental design and analysis issues related to timing based on three exploratory studies investigating imitation based interaction games with child-sized humanoid robots and human participants. The primary goal of this work is to achieve (non-verbal) gesture communication and imitation between child-like humanoid robots and human beings, whereby interaction games including drumming and imitation served as a test bed to study key aspects of face-to-face interaction such as turn-taking, synchronisation and non-verbal gestures.

The first presented study is based on *drum-mate*, a drumming game where turn-taking is deterministic and head gestures of the child-sized robot KASPAR[1] accompany its drumming to assess the impact of non-verbal gestures on the interaction [2]. This paper will focus on a modified version based on emergent turn-taking dynamics; here our aim is to have turn-taking which is not deterministic but emerging from the social interaction between the human and the humanoid [2].

The second study focused on imitation. Unlike the first work, which concerns turn-taking, this work is based on synchronisation which introduces different issues related to timing. Here the robot makes simple body moves like waving its hand, and the human tries to imitate the robot whilst the robot evaluates how successful the imitation is [3]. Here the joint motion of the human and the robot should be tracked simultaneously and compared.

In a third study, our aim is to analyse and model the gaze behaviour of human-human and human-humanoid pairs. Therefore we need to track the gaze of the participants coming from different sources in real-time and compare them to detect joint and mutual gaze. Additionally, once the data is collected, a suitable representation for the time distribution of the periods of mutual gaze must be chosen.

3. Issues related to timing

A. Turn-taking issues

We implemented the human-robot drumming game as an example of a *call and response turn-taking interaction.* In the deterministic case, we used predefined fixed time duration heuristics for turn-taking . The human partner started by playing simple rhythms with a toy drum. KASPAR started playing if the human was silent for a few seconds. However, it was not always clear when the robot or human partner should initiate interaction in taking a turn. In the second version of the study (the emergent case), we instead used probability-based computational models to control timing and turn-taking. Three simple models (threshold, linear, hyperbolic) were used to control the starting and stopping of the robot's drumming beats. The temporal dynamics of turn-taking thus emerged from the interaction between the human and the humanoid. We studied how these models impacted the drumming performance of the human-robot pair and the participants' subjective evaluation of the drumming experience.

B. Synchronisation issues

Synchronisation is another vital issue in timing which we encounter during human-robot interaction

[1] Acknowledgements: This work was conducted within the EU Integrated Project RobotCub ("Robotic Open-architecture Technology for Cognition, Understanding, and Behaviours"), funded by the EC through the E5 Unit (Cognition) of FP6-IST under Contract FP6-004370.

experiments. Many interaction games, including physical imitation or music, are based on synchronisation. The second experiment we present was based on the synchronisation of the human and humanoid using a simple arm waving motion. We used magnetic motion trackers to detect humans' arm motion and compared with the position of the robot's arm joints, and tried to detect the synchronisation between them in different scenarios, i.e. waving hands totally in phase or out of phase [3]. We proposed a method based on *information distance* to detect the similarity and synchronisation between the motion of human and the humanoid robot KASPAR2 [3,4].

Fig. 1 A screen shot from the experiments

C. Issues related to measurement devices

In the third study, we analysed the gaze habits of two human subjects. Human-human pairs sat at each side of a table looking towards each other and had simple dialogues while their gazes were tracked by special eye-tracker devices to analyse their joint gaze habits.

Timing is a big issue in using cameras, eye-trackers and other devices to measure body motion, which are essential for human-robot interaction studies. The data coming from these devices should be time-stamped to use with data from other sources. If data is collected on different PCs/laptops, synchronizing time-stamps is not trivial. A Network Time Protocol (NTP) server/client setup should be able to maintain clock accuracy among machines within tens of milliseconds, a resolution which should be adequate for most sensors [5]. Better synchronization is commonly reported, especially across local networks, but may not be reliably achievable without specialized hardware or software [6].

D. Issues related to adaptation:

Adaptive behaviour is a very important part of our interactive studies and in the case of producing or detecting such behaviour, timing is very crucial. It is important to compare the real-time waving-motion data/ drumming performance/gaze direction of the robot and the human (or human-human pair) to get feedback which will be used in the adaptation. If the data from both participants can not be synchronized correctly this feedback can not be achieved.

Additionally, creating adaptive behaviour often relies upon designing or learning a computational model of the desired behaviour. This model may be, as in the case of the drumming study, quite simple, but the realism and interpretability of the behavior produced is likely to be highly dependent upon its internal representation. Because of the importance of timing in the interactions we've explored, models of these behaviours must explicitly represent timing relationships in order to capture their fundamental characteristics. What characteristics of timing (e.g, duration, periodicity, tempo) should be represented and how best to represent them (as values, distributions, or functions) is highly dependent on the nature of the interaction, and many have to be determined via trial-and-error or by examining data from humans performing the behaviour in question.

3. Conclusion

We presented the experimental design and the timing related issues resulting from three interaction studies with child-sized humanoid robots and human participants. Timing plays an important role in human-humanoid interaction, appearing in several different ways such as turn-taking, synchronization, real-time interaction and adaptation.

The methodologies and solutions to the related issues presented in this paper will be used for future studies related to social interaction between human and humanoid, and can possibly be extended for use in other robotic fields, e.g. entertainment, service, and educational/therapy robots.

Selected References

[1] Blow, M.P, K. Dautenhahn, A. Appleby, C. Nehaniv, D. Lee, "Perception of robot smiles and dimensions for human-robot interaction design*", In Proc. IEEE RO-MAN06,* pp.469-474, 2006.

[2] Kose-Bagci, H., K. Dautenhahn, D. S. Syrdal, and C. L. Nehaniv, "Drum-mate: Interaction dynamics and gestures in human-humanoid drumming experiments", Connection Science, 2009, accepted.

[3] Shen, Q., J. Saunders, H. Kose-Bagci, K. Dautenhahn, "Acting and Interacting like me? A method for identifying similarity and synchronous behaviour between a human and robot", Proceedings IEEE IROS workshop on "From motor to interaction learning in robots", Nice, France. 2008

[4] Shen, Q., J. Saunders, H. Kose-Bagci, K. Dautenhahn, "An Experimental Investigation of Interference Effects in Human-Humanoid Interaction Games", IEEE RO-MAN2009,, accepted

[5]Mills, D. L. 1989. On the Accuracy and Stablility of Clocks Synchronized by the Network Time Protocol in the Internet System. *SIGCOMM Comput. Commun. Rev.* 20, 1 (Dec. 1989), 65-75.

[6]Pásztor, A. and D. Veitch, PC based precision timing without GPS. In *Proceedings of the 2002 ACM SIGMETRICS* June 15 - 19, Marina Del Rey, California .2002.

Connection Science
Vol. 00, No. 00, Month 2009, 1–32

Taylor & Francis
Taylor & Francis Group

Copyedited by: KS

Drum-mate: interaction dynamics and gestures in human–humanoid drumming experiments

Hatice Kose-Bagci*, Kerstin Dautenhahn, Dag S. Syrdal and Chrystopher L. Nehaniv

Adaptive Systems Research Group, School of Computer Science, University of Hertfordshire, Hatfield, UK

(*Received 25 July 2008; final version received 29 September 2009*)

This article investigates the role of interaction kinesics in human–robot interaction (HRI). We adopted a bottom-up, synthetic approach towards interactive competencies in robots using simple, minimal computational models underlying the robot's interaction dynamics. We present two empirical, exploratory studies investigating a drumming experience with a humanoid robot (KASPAR) and a human. In the first experiment, the turn-taking behaviour of the humanoid is deterministic and the non-verbal gestures of the robot accompany its drumming to assess the impact of non-verbal gestures on the interaction. The second experiment studies a computational framework that facilitates emergent turn-taking dynamics, whereby the particular dynamics of turn-taking emerge from the social interaction between the human and the humanoid. The results from the HRI experiments are presented and analysed qualitatively (in terms of the participants' subjective experiences) and quantitatively (concerning the drumming performance of the human–robot pair). The results point out a trade-off between the subjective evaluation of the drumming experience from the perspective of the participants and the objective evaluation of the drumming performance. A certain number of gestures was preferred as a motivational factor in the interaction. The participants preferred the models underlying the robot's turn-taking which enable the robot and human to interact more and provide turn-taking closer to 'natural' human–human conversations, despite differences in objective measures of drumming behaviour. The results are consistent with the temporal behaviour matching hypothesis previously proposed in the literature which concerns the effect that the participants adapt their own interaction dynamics to the robot's.

Keywords: social robots; humanoids; robot drumming; human–robot interaction; interaction kinesics; emergent turn-taking

1. Introduction

The development of socially intelligent and adaptive robots in human–robot interaction (HRI) is an emerging interdisciplinary field across the boundaries of robotics, engineering and computer science on the one hand, and psychology, ethology and social sciences on the other (Dautenhahn 2007a). The primary goal of our research is to design a 'successful' HRI, whereby the robot is engaged in certain tasks and carries out these tasks in a manner that is socially appropriate, for example, enjoyable and acceptable for its users (Dautenhahn 2007b). It remains an open research challenge to design such 'successful' HRI: success is here defined in terms of both performance

*Corresponding author. Email: h.kose-bagci@herts.ac.uk

ISSN 0954-0091 print/ISSN 1360-0494 online

DOI: 10.1080/09540090903383189
http://www.informaworld.com

of the human–robot pair in a task-based scenario, as well as in terms of the user's subjective experience of the interaction. Intuitively one may assume that what matters in human–human interaction should also matter in human–machine interaction. And indeed, research by Nass and his colleagues (e.g. Reeves and Nass 1997; Nass and Lee 2000) has shown that people treat interactive artefacts socially. However, robots and computers are not exactly like people and it remains open when and to what extent models and theories of human–human interaction are directly applicable to HRI (Dautenhahn 2007b).

In this article, we are particularly concerned with the *dynamics* of HRI. Specifically, we address the question of whether details of the dynamics of interaction that have been shown to play a fundamental role in human–human interaction are equally important in HRI. In human–human interaction, details of timing and synchronisation of gestures, speech, turn-taking in interaction, etc. influence the nature and meaning of interaction. But is the same also true of HRI? Implementing sophisticated dialogue and interaction models between humans and machines requires significant computational and research effort. In order to decide whether this effort is justified, we need to demonstrate that details of HRI kinesics matter. To address this issue, we used in our experiments simple and (algorithmically) arbitrary, minimal computational models underlying the robot's turn-taking dynamics, rather than trying to model faithfully complex mechanisms of cognition and learning in humans. We argue that if our simple models show an effect, that is, if we find that the details of simple interaction dynamics significantly influence the 'success' of the interaction (both in terms of objective performance and subjective user evaluation), then these results suggest that future research in HRI design needs to take into account the details of robot interaction dynamics even when not strictly based on cognitively plausible models of turn-taking and interaction.

The work discussed in this article is related to our wider research agenda where we study the importance of timing, rhythms, turn-taking and entrainment, which are key factors in the development of communication (cf. Robins et al. 2005; Robins, Dautenhahn, te Boekhorst, and Nehaniv 2008). Communication is an integral part of human social interaction. Developmental psychologists distinguish between: (a) a primary, expressive system which has semantic and intentional content but does not take account of the communication partner,[1] and (b) a pragmatic, referential system which can predict, and infer intention in the communication partner (Nadel, Guerini, Peze, and Rivet 1999). These two key processes are involved in supporting a transition from primary to pragmatic communication which requires mastering interpersonal timing and the ability to communicate about a shared topic. Research has identified the importance of contingency in rhythm, timing and inter-subjectivity in early communicative interaction of infants with a caregiver. Such protoconversation plays a key role in the natural developmental progression of human infants (Trevarthen 1999). Detailed analyses of infant–caretaker interactions show that turn-taking between adult and infant in these protoconversations are closely coordinated and reach rapid mutual entrainment.

Even before the link has been made to infant development, researchers studying human–human interaction had long recognised the importance of timing, turn-taking and synchronisation dynamics (Condon and Ogston 1967; Kendon 1970; Hall 1983). Goldin-Meadow argues that the gestures the people produce in their conversation are tightly intertwined in their timing and meaning, and that non-verbal gestural components of people's communication cannot be separated from the content of conversation (Goldin-Meadow and Wagner 2005). According to Bernieri and Rosenthal, '[i]nterpersonal coordination is present in nearly all aspects of our social lives, helping us to negotiate our daily face-to-face encounters… We also coordinate our non-verbal behavior with others to communicate that we are listening to them and want to hear more' (Bernieri and Rosenthal 1991, p. 401). In this context, interpersonal coordination is loosely defined as '…the degree to which the behaviors in an interaction are nonrandom, patterned, or synchronised in both timing and form' (Bernieri and Rosenthal 1991, p. 403).

Within the wider context of interpersonal coordination, in our work we focus on *interaction kinesics*, which can be described as the study of the role and timing of non-verbal behaviour, including body movements, in communicative and interactional dynamics. While numerous studies have investigated how people adapt to other humans (e.g. Pickering and Garrod 2004), non-human stimuli (e.g. Schmidt, Richardson, Arsenault, and Galantucci 2007) or computers (e.g. Suzuki and Katagiri 2007), interaction kinesics in HRI is a relatively unexplored area of research (Robins et al. 2005, 2008). And only few studies have focussed on experimental investigations of this important topic. For example, Watanabe (2004) investigated the embodied entrainment between speech and body motions such as nodding in face-to-face communication involving robotic and virtual characters engaging with people. Yoshikawa, Shinozawa, Ishiguro, Hagita, and Miyamoto (2006) highlighted the role of responsive gaze in human–humanoid interaction. Yamamoto and Watanabe (2003) found the differences in people's preferences concerning the timing of utterances in human–robot greeting interactions. Robins et al. (2008) explored interaction kinesics in child–robot interaction in a play context involving a robotic dog (Reeves and Nass 1997) and the child-sized humanoid KASPAR.[2] Yamaoka, Kanda, Ishiguro, and Hagita (2007) showed in an experiment with the Robovie robot and student participants how the contingency of interaction impacts participants' perception of the autonomy of the robot, depending on the degree of complexity of the interaction. The role of Robovie's response time as well as strategies of how a robot can cope with delays has been investigated by Shiwa, Kanda, Imai, Ishiguro, and Hagita (2008). A recent study by Yamaoka, Kanda, Ishiguro, and Hagita (2008) with Robovie studies the effect of the robot's body position and orientation on people's proxemics behaviour in joint attention scenarios. Outside the context of interactive robots, the importance of timing and synchronisation has also been studied in human–computer interaction (Suzuki and Katagiri 2007) and has been applied to therapeutic walking devices (Miyake 2003), as well as in evolved artificial social turn-taking agents (Iizuka and Ikegami 2004). The earlier-mentioned examples indicate the growing interest of the HRI community in interaction kinesics.

The particular experimental context chosen in our work is that of human–robot drumming. We decided to choose a joint drumming task since collaborative music performance, in general, lends itself to the study of interaction between humans and robots involving a variety of social aspects including imitation, gestures, turn-taking and synchronisation, occurring in an overall playful and enjoyable context. From a robotics point of view, drumming is a very suitable means of performing music, since it is relatively straightforward to implement and test, and can be realised technically without special actuators like fingers or special skills or abilities specific to drumming. Thus, the drumming scenario provides a playful and interactive context that allows to constrain and manipulate different experimental parameters easily.

Several researchers have studied drumming in the context of human–robot music performance. In Weinberg, Driscoll, and Parry (2005), Weinberg and Driscoll (2006) and Crick, Munz, and Scassellati (2006), robotic percussionists play drums in collaboration with interaction partners. In Weinberg et al. (2005), an approach based on movement generation using dynamical systems was tested on a Hoap-2 humanoid robot using drumming as a test case. Similarly, in Kotosaka and Schaal (2001), humanoid drumming is used as a test bed for exploring synchronisation. However, none of the prior work has specifically studied the socially interactive aspects in general, or interaction kinesics in particular, in the context of human–humanoid drumming, which are the focus of this article.

In this article, we present the results from two empirical studies involving adult participants[3] interacting with the humanoid robot KASPAR in an imitation-based interaction game based on drumming. The two experiments highlight the different aspects of HRI: (a) the role of (non-verbal) gesture communication in a joint drumming task, and (b) the dynamics of emergent turn-taking games. In Section 1.1, we will motivate the first experiment based on gesture communication which used non-verbal gestures as social cues. This approach is discussed in the light of related work on several robotic percussionists, as well as other work in the wider context of social robotics. In Section 1.2, we motivate our work on emergent dynamics of turn-taking interaction, in the context of literature highlighting the importance of turn-taking in conversations and interaction games. The actual experiments will be described in Sections 2 and 3. Note that the field of social robotics and HRI is very active, with a variety of different robotic systems used in interaction studies. A complete review of the literature in this field goes beyond the scope of this experimental paper; so we will focus our discussion of related work on research specifically relevant to our research questions. For a very recent review of the field of HRI, see Goodrich and Schultz (2007).

1.1. *Gesture communication: motivation and related work*

A robot that engages with people in interaction games could benefit from behaviour that specifically motivates the user and sustains the interaction while coping with a wide range of users. One way of motivating people to interact is through the use of social cues such as gestures. In human–human interaction, gestures play an important role in communication, coordination and regulation of joint activities. Indeed, in the related field of virtual agents, researchers have shown the beneficial effects of gestures and expressions used by virtual agents, both in short-term and in long-term interactions, in maintaining user involvement with the tasks encouraged by the agent (Bickmore and Cassell 2005; Bickmore and Picard 2005).

Applied to robotics, this suggests that a robot may require social cues and gestures to motivate users to interact with it, for example, in the field of assistive robotics (Tapus and Matarić 2006). A variety of robotic systems have been using social cues and gestures to encourage HRI. A well-known example is KISMET, where facial expressions were used to regulate the interaction with people inspired by interactions of infants with their caretakers (Breazeal 2002). Other recent examples include small cartoon-like robotic ‘creatures’ such as KEEPON and ROILLO, designed to be used in interaction with children (Kozima, Nakagawa, Yasuda, and Kosugi 2004; Michalowski, Sabanovic, and Michel 2006). These small robots have a limited action repertoire, but can produce selected gestures to engage in interaction with children in the playground. The fixed gestures are either random or tele-operated by a hidden puppeteer via a Wizard of Oz technique, as a part of social interaction. ROILLO is a simple robot with a rubber coated foam head, body and an antenna. It has three wires connected to simple servos, which move the head and body in various directions. It is used in experiments to study the interactions between the robot and the children (Michalowski et al. 2006). KEEPON is a minimalist expressive robot that only has a rubber head and an oval body. It has a small CCD camera and a microphone on it. It can move its head, turn its body and make bobbing actions to show its ‘feelings’. It has both attentive and emotive actions. It is simple but robust enough to be used in play rooms in interaction with children (Kozima et al. 2004; KEEPON 2007, http://univ.nict.go.jp/people/xkozima/infanoid/robot-eng.html#keepon).

Related work on human–robot drumming includes HAILE (Weinberg et al. 2005; Weinberg and Driscoll 2006), a robot arm designed specifically to drum in dynamic and musically sophisticated collaboration with creative human musicians. HAILE does not use fixed deterministic rules, but uses autonomous methods to create variant rhythms. It perceives a variety of complex features of the human partner’s drumming, analyses the sound patterns and produces rhythms in response. Compared with HAILE, in Crick et al. (2006) a less musically sophisticated humanoid robot called NICO with an upper half body torso plays a drum together with human drummers. It has visual and audio sensing to determine an appropriate tempo adaptively using a simple threshold mechanism to parse the human partner’s beats, and can distinguish its own performance with audio sensing, integrating the two sources of information to predict when to perform the next beat.

The above motivation and background led to our first experiment, where the humanoid robot KASPAR plays the drums autonomously with a human 'partner' (interactant), trying to imitate the rhythms produced by the human while using non-verbal gestures to motivate the human. In this experiment, KASPAR's behaviour is deterministic in the sense of producing the same (actuator) output given the same input from its sensors.[4] KASPAR produces non-verbal (head) gestures from a limited repertoire and eye-blinking as it drums. Our approach is tested using different degrees of such non-verbal gesturing with adult participants in several drumming sessions, and the experimental results are reported and analysed below (Section 2) in terms of imitation, turn-taking and the impact of non-verbal gestures as social cues.[5]

1.2. *Emergent turn-taking dynamics: motivation and related work*

Turn-taking is an important ingredient of human–human interaction and communication, whereby the role switch ('leader' and 'follower') is not determined by external sources but emerges from the interaction. Human beings generally 'know' when to start and stop their turns in the social interactions, based on various factors including the context and purpose of the interaction, feedback from the social interaction partners, emotional and motivational factors, etc. They use different criteria for these decisions. In this work, our aim is to build a framework which enables emergent turn-taking, and role-switching between a human and a humanoid in an imitation game, and to understand how differences in robot turn-taking strategy can influence the emergent dynamics of HRI. We do not aim to produce psychologically plausible models of human turn-taking behaviour in this work, but employ simple, minimal generative mechanisms to create different robotic turn-taking responses/strategies.

Related work that studied turn-taking in games and conversations focussed on different aspects. An example from developmental psychology is described in Hendriks-Jansen (1996), which discusses emergent turn-taking between a mother and a baby without any explicit 'control' mechanism (e.g. the mother starts jiggling in response to her baby's sucking to encourage her baby to resume sucking). This results in emergent turn-taking between the jiggling and the sucking actions. Turn-taking also has important implications in robot-assisted therapy. Indeed, one therapeutically relevant issue in teaching and education of children with autism is to teach children the concept of 'turn-taking'. Turn-taking games have been used to engage children with autism in social interactions (Dautenhahn and Billard 2002; Robins, Dautenhahn, te Boekhorst, and Billard 2004a).

Another example of turn-taking games is given from a cognitive robotics view in R.A. Brooks (personal communication, August 28, 1997). In this work, a ball game between a humanoid robot COG and the human experimenter is described. COG and the human were reaching out and grasping a ball in alternation. Note that in this case the experimenter led the turn-taking behaviour in reaction to the robot's visually driven actions. Ito and Tani (2004) studied joint attention and turn-taking in an imitation game played with the humanoid robot QRIO, where the human participants tried to find the action patterns, which were learned by QRIO previously, by moving synchronously with the robot.

From a linguistics point of view, some of the important features of turn-taking in human conversation identified are as follows (Sacks, Schegloff, and Jefferson 1974):

- Speaker-change recurs, or at least occurs.
- Mostly, one party talks at a time.
- Occurrences of more than one party speaking at the same time are common but brief.
- Transitions (from one turn to the next) with no gap and no overlap are common (slight gap or slight overlap is accepted).
- Turn order is not fixed, but varies.

- Turn size is not fixed, but varies.
- Length of conversation is not specified in advance.
- What parties say is not specified in advance.
- Relative distribution of turns is not specified in advance.
- Number of parties can vary.
- Talk can be continuous or discontinuous.

Built on these features, Thórisson (2002) developed a turn-taking mechanism for conversations based on his previous work on the so-called Ymir mind model for communicative creatures and humanoids. He proposed, implemented and tested a generative, multi-modal turn-taking model for a face-to-face dialogue. The model was based on literature in human–human dialogue. The above-mentioned expressive humanoid robot KISMET (Breazeal 2002, 2003) which used social cues for regulating turn-taking in non-verbal interactions with people used a sophisticated robot control architecture modelling motivation, emotions and drives to satisfy KISMET's internal 'needs'. Turn-taking between KISMET and humans emerged from the robot's internal needs and goals and its perceptions of cues from its interaction partner. Rather than trying to model any particular turn-taking behaviour as observed in human–human dialogue (as it has been done e.g. in Thórisson's (2002) work mentioned above), we pursued a synthetic, bottom-up approach by defining very simple models of turn-taking based on basic mathematical functions. Such a bottom-up approach is in line with other approaches in the research field of Embodied Artificial Intelligence (Steels and Brooks 1995; Pfeifer and Scheier 1999) and is here applied to human–humanoid interaction aiming at developing socially interactive behaviour for a humanoid robot.

Also, different from the above-mentioned work with KISMET, where the interaction was the goal in itself, we wanted to include a certain (enjoyable) task that needs to be achieved jointly by the human–robot pair, to provide the overall context.

Important in this context is the temporal behaviour matching hypothesis as proposed in Robins et al. (2008), which predicts that in HRI games, people will adapt to and match the robot's temporal behaviour, similar to the effects that can be found in the literature of human–human interaction. The hypothesis has been supported in experiments with children who were playing imitation games with KASPAR (the same robot as used in our experiments; Robins et al. 2008). While this hypothesis may at first seem trivial since people and other animals are very adaptive and adapt to the dynamics of a variety of stimuli (see, e.g. Schmidt et al. 2007), for roboticists it is very important to actually know whether people do indeed adapt and respond to the dynamics of robot behaviour – if it were false then one would not need to take robot interaction dynamics and kinesics into account – which would substantially simplify HRI design. Moreover, what types of impact robot kinesics can have on interaction and the degree and manner in which different people might be influenced differently are open issues. Thus, for HRI researchers, this is an important question to study experimentally, and, as discussed in more detail above, it has only recently attracted attention in the field of robotics and HRI (c.f. Robins et al. 2005, 2008; Crick et al. 2006; Yoshikawa et al. 2006).

Based on the above motivation and background, we designed a second experiment where KASPAR plays the drums autonomously with a human 'partner' (interactant), trying to imitate the rhythms produced by the human (as a follower) and trying to motivate (as a leader in the game) the human to respond. Using different simple, probabilistic models, KASPAR decides when to start and stop its turn. It observes the human playing and uses its observations as parameters to decide whether to listen to the human or to take the turn actively in the game. This is different from Experiment I where we tested deterministic turn-taking. This work was tested with adult participants and the results were studied in terms of imitation, interaction and turn-taking.[6]

The two experiments are described below in detail separately due to clear differences in research questions and implementation of the interaction games. However, both experiments share a common methodological approach.

We chose a within-participant design for both studies for two main reasons: (a) the study of individual differences as such is an interesting challenge in HRI research (Breazeal 2004) and (b) previous research has indeed found significant individual differences in HRI studies, for example, concerning personality traits (Walters, Syrdal, Dautenhahn, te Boekhorst, and Koay 2008), gender and personality (Syrdal, Koay, Walters, and Dautenhahn 2007), human and robot personality matching (Tapus, Tapus, and Mataric 2008), and user personality and robot personality style (Wrede, Buschkaemper, and Li). Since the literature shows individual differences of how people respond in HRI studies (e.g. based on the participants' gender, age, individual personality traits, etc.), a within-participant design approach thus seemed most suitable for understanding the range and variability, and impact of robot kinesics on interactions.

In both experiments, we evaluate the objectively measured performance of the human–robot pair as well as the subjective interaction experience as judged by the human participants.

The rest of this article is organised as follows. In Section, the first experiment on deterministic turn-taking is presented, followed by Section 3, which describes the second experiment on emergent dynamics of turn-taking. Each of these two experimental sections includes the corresponding research questions as well as descriptions of the experimental setup, experimental results and discussions of the results. Section 4 presents the overall conclusion. The final section of this article outlines the ideas for future work.

2. Experiment I: deterministic turn-taking

2.1. *Methodology*

In the first experiment, the human partner played a rhythm which KASPAR tried to replicate, in a simple form of imitation (mirroring). KASPAR has two modes: listening and playing. In the listening mode, it recorded and analysed the played rhythm, and in the playing mode, it played the rhythm back by hitting the drum positioned in its lap. Then the human partner played again. This (deterministic) turn-taking continued for the fixed duration of the game. KASPAR did not imitate the strength of the beats but only the number of beats and duration between beats. For beat frequencies beyond its skill, it used instead minimum values allowed by its capabilities.[7] It also needed a few seconds before playing any rhythm to get its joints into correct reference positions.

Figure 1 presents the basic model of KASPAR–human interaction. The model requires the gestures of both human and humanoid for social interaction, as well as drumming. Human gestures or body movements were not detected in our experimental setup and were therefore not considered in the implementation.

One of the fundamental problems in this scenario is the timing of the interaction; as discussed above, timing plays a fundamental role in the regulation of interaction. It is not always clear when the robot or human partner should start interaction in taking a turn. In this experiment, the model used some predefined fixed time duration heuristics for synchronisation. KASPAR started playing if the human partner was silent for a few seconds, and tried to motivate the participant with simple gestures.

2.2. *Research questions and expectations*

Our primary research question concerned the possible impact of robot gestures on the imitation and turn-taking game (in terms of performance), but also on the participant's subsequent evaluation of

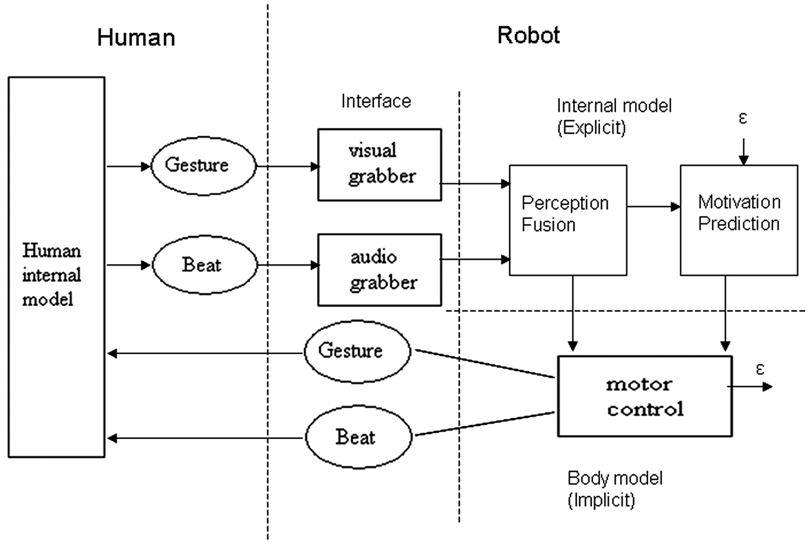

Figure 1. The model for KASPAR–human interaction.

the game. We expected that the participants would be more engaged and would evaluate the interactions more positively in experimental conditions where KASPAR used non-verbal head gestures. Moreover, we expected that too many gestures may distract people from the drumming task.

2.3. *Experimental conditions*

We studied three conditions with increasing amounts of robot gesturing:

(1) No-gesture: KASPAR did not use any gestures, it only imitated the human drumming beats it detected.
(2) Gesture: Simple head gestures (e.g. moving the head to the right or left, moving the head up or down, tilting the head slightly to different angles) and eye blinking were included in KASPAR's movements. KASPAR started drumming using one of a fixed set of gestures. If the human partners did not play their turn, then KASPAR did not respond either, and then the turn passed back to the partner. A fixed order of *n* gestures was used, and this order was repeated for every *n* turns. It was intended that the value for *n* should be large enough so that the participant would not realise that this was a fixed pattern but rather that the gestures seem either 'meaningful' or random (in the experiment, *n* was set to seven based on simulated experiments, i.e. carried out with the experimenter as the interaction partner).
(3) Gesture+: This condition is the same as *gesture*, except that KASPAR displayed on its turn in the interaction gestures even when neither the robot nor the participant played the drum. The gestures used were the same as in the *gesture* condition, and the drumming part was the same in all the three conditions.

2.4. *Experiment, results and analysis*

2.4.1. *Robot*

The experiment was carried out with the humanoid robot called KASPAR (Figure 2). KASPAR is a humanoid robot that has been designed specifically for HRI studies. It possesses a minimal set

Figure 2. The humanoid robot KASPAR and its toy drum that were used in the experiments.

of expressive robot features (cf. Blow, Dautenhahn, Appleby, Nehaniv, and Lee 2006) for more information on its design rationale. KASPAR has eight degrees of freedom (DOF) in the head and neck, and six in the arms and hands. The face is a silicon-rubber mask, which is supported on an aluminium frame. It has 2 DOF eyes fitted with video cameras, eye lids that allow blinking and a mouth capable of opening and smiling; see Blow et al. (2006) for a more detailed description.

2.4.2. *Experimental setup*

The experiment was carried out in a separate room isolated from other people and noises which could affect the drumming interaction. KASPAR was seated on a table with the drum positioned on its lap. The participants were seated in front of the robot using another drum that was fixed on the table (Figure 3). The participants used a pencil to hit the drum. Although we suggested to

Figure 3. A video snap shot from the experiments.

the participants to use one pencil and hit on the top of the drum, sometimes they used two pencils with a single hand or with both hands, and several times they used the tambourine-style bells around the drum's sides.

2.4.3. *Software features*

The implementation of robot perception and motor control used the YARP environment (Metta, Fitzpatrick, and Natale 2006). YARP is an open-source framework that supports distributed computation that emphasises robot control and efficiency. It enables the development of software for robots, without considering a specific hardware or software environment. Portaudio (2007; http://www.portaudio.com/trac/wiki/) software was used to grab the audio from the audio device, within the YARP framework. See Appendix 1 for details of the audio analysis.

2.4.4. *Participants*

Twenty-four participants (7 female and 17 male) took part in the study. Due to logistical reasons, the trials were carried out in 2 sets (a few months apart) with 12 participants each. All the participants worked in computer science or similar disciplines at the university. Only six of them had interacted with KASPAR prior to the experiment, and most of the participants were not familiar with robots in general. Note that we initially did not plan to study the influence of gender in the experiment; for this reason, the sample is not gender-balanced. However, where appropriate we mention gender differences that were observed. Four of our participants had children.

2.4.5. *Interaction game setup*

We used a 1 min demonstration of the robot without any drumming game play, where the participants were shown how to interact with KASPAR. This was followed by three games reflecting the three experimental conditions described above each lasting 3 min, without pointing out to the participants any differences between the conditions. We presented the game conditions in all the possible six different orders to analyse the effect of the order of the games. To account for possible fatigue or habituation, in the sequential order section below, we analysed the games according to their order number in the sequence experienced by the participants (independent of the particular experimental condition), as being the first game, second or third, disregarding their game types, for example, for one participant the first game (number 1) would be the no-gesture game, and for another participant, no-gesture would be the third game (number 3). After each participant finished the three games, they were asked to complete a questionnaire to assess how they subjectively evaluated the three different games.

2.4.6. *Results*

2.4.6.1. Evaluation of questionnaire data. The participants were invited to evaluate their interaction with KASPAR using a questionnaire. There were two items inviting the participant to choose which game was the most and least preferred overall. There were also three five-point Likert scales which allowed the participant to rate each drumming game in terms of (1) how much they enjoyed the game, (2) how well KASPAR drummed and (3) how sociable they perceived KASPAR to be. Open-ended questions were included to allow participants to explain their reasoning for their preferences. Most and least preferred games according to game types and sequential order were statistically analysed using a χ^2 test.

Most and least preferred games according to game type: The frequencies of participants which rated each game as most preferred and least preferred are presented in Table 1 along with residuals based on an expected count of 7.7. The differences from the expected counts were significant for both the most preferred game type ($\chi^2(2) = 6.61$, $p = 0.037$) and the least preferred game type ($\chi^2(2) = 9.74$, $p = 0.008$). The majority of the participants preferred the *gesture* game and disliked the *no-gesture* game. Their general opinion was that the game without gestures was also poor in terms of social interaction and enjoyment, which encouraged them to play more. For the *gesture* game, they said they prefer the right balance of drumming and interaction.

Most and least preferred games according to sequential order: A significant difference was found between the first and third games in terms of sequential order ($\chi^2(1) = 4.57$, $p = 0.033$). There is no significant difference overall between the three games if the second game is included (Table 2). Open-ended responses highlighted that the majority would become more familiar with the game as they played more, allowing them to interact more efficiently with KASPAR in terms of the drumming tasks. Another issue raised in the open-ended responses was that the participants would become fatigued and bored after doing the repetitive drumming task for a prolonged period of time, which may explain the lack of a significant difference between the second and third games.

Preferences: While the method of counterbalancing is an accepted means of protecting against confounders due to presentation order (Miller 1984), the clear main effect of presentation order was considered a threat to this assumption. To control for this threat, mixed model ANOVAs were run using game type to investigate possible interaction effects of presentation order and game type on both questionnaire responses and behavioural data. These were mainly non-significant, supporting the notion of independence between presentation order and game-type overall in the sample. The one exception is addressed in Section 2.4.6.2.

Sample similarities: In terms of differences between the first sample of 12 and the second sample of 12 participants, a mixed-model ANOVA found no significant differences in terms of preferences ($F(1,22) = 0.772$, $p = 0.39$). Thus, in the following we present the results from the overall sample of 24 participants.

Table 1. Most and least preferred games according to game types.

	Participants			
Game Type	Most preferred	Residual	Least preferred	Residual
No-gesture	3	−4.7	12	5.3
Gesture	12	5.3	1	−6.7
gesture+	7	−0.7	9	1.3
No preference	2	N/A	2	N/A

Table 2. Most and least preferred games according to sequential order.

	Participants			
Order	Most preferred	Residual	Least preferred	Residual
1	3	−4.3	10	−4.3
2	8	5.3	5	−0.7
3	11	−0.7	8	3.7
No preference	2	N/A	1	N/A

Preferences according to sequential order: The preferences according to order within the sample as a whole were assessed using a repeated measures ANOVA. There was an effect approaching significance in how participants rated KASPAR's drumming according to the order of the game ($F(2,46) = 3.11$, $p = 0.054$). No significant effects on game order were found in terms of the robot's sociality or enjoyment ratings. Participants tended to rate the last game more favourably across the different rating types (despite the fatigue reported by some participants during later games), see Figure 4. The results from the ANOVA, as well as the descriptives described in Figure 4, suggest that this trend was the most pronounced in the way the participants rated KASPAR's drumming.

Preferences according to game type: The repeated measures ANOVA for preferences dependent on game type found an effect approaching significance in terms of how KASPAR's drumming was rated according to game type ($F(2,46) = 2.71$, $p = 0.077$) as well as for general enjoyment of the game ($F(2,46) = 2.81$, $p = 0.07$). We found a significant effect for game type in terms of how KASPAR's sociality was rated ($F(2,46) = 5.01$, $p = 0.011$), see Figure 5.

Figure 5 suggests different trends for the different game types. The trend approaching significance for KASPAR's drumming suggests that the drumming aspect of the interaction for the *no-gesture* game was rated the most favourable, followed by the *gesture* game, with the *gesture+* game receiving the lowest rating.

In terms of the social aspect of the interaction, the opposite effect was found. The *no-gesture* game was rated the lowest, with the *gesture* and *gesture+* games rated higher. For overall enjoyment, the *gesture* games were rated the highest, followed by *gesture+*. The *no-gesture* game was rated the lowest.

2.4.6.2. Evaluation of behavioural data. The behavioural data required for the evaluation of the participant's and the robot's performance during the games were collected based on the data on the robot's own drumming behaviour and video recordings of the human's drumming behaviour which were annotated manually and then analysed quantitatively. The behavioural data include the number of turns in a specific game, the number of drumming bouts performed by the participants and the robot, and the 'drumming errors'. The errors are the differences between KASPAR's

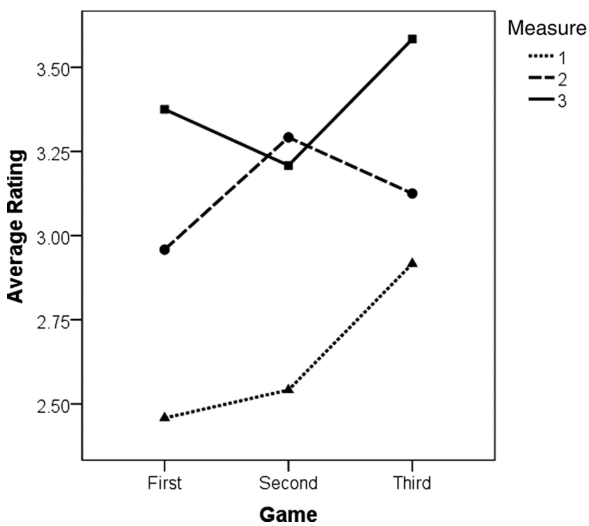

Figure 4. Ratings for games according to order in terms of (1) KASPAR's drumming, (2) KASPAR's sociality and (3) enjoyment of the game.

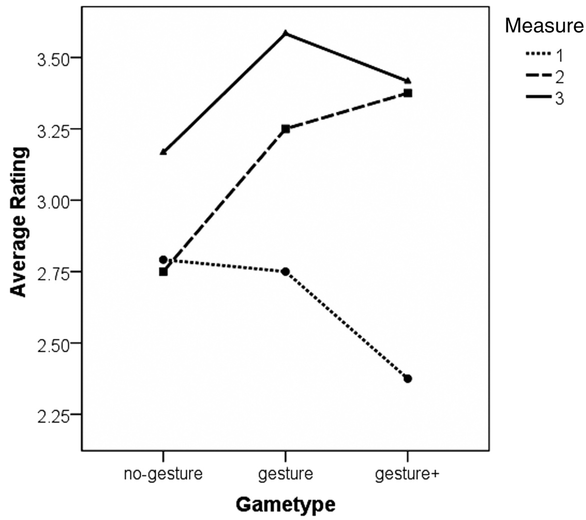

Figure 5. Ratings for games according to game type in terms of (1) KASPAR's drumming, (2) KASPAR's sociality and (3) enjoyment of the game.

actual drumming (i.e. the number of beats KASPAR plays in a particular turn) and the number of beats the participant plays. We calculated an average error per turn. Thus, 'errors' do not reflect any mistakes in the system as such, but reflect the discrepancy between human's and robot's drumming performance.

Behavioural data according to sequential order: We found a significant effect for sequential order in terms of average number of errors ($F(2,46) = 6.18$, $p = 0.004$). This effect is seen in Figure 6 and suggests that the errors were in general lower for later games.

Generally, the participants either tried very long and fast patterns or they did not beat loud enough to be detected reliably (KASPAR uses a high-level noise filter to filter out high inner noise coming from its joints, so it can only sense loud beats) when they started to play initially. Interestingly, without any external encouragement, as they got used to the game, they progressively synchronised their drumming to the robot. Details of the results are presented in Table 3. As such,

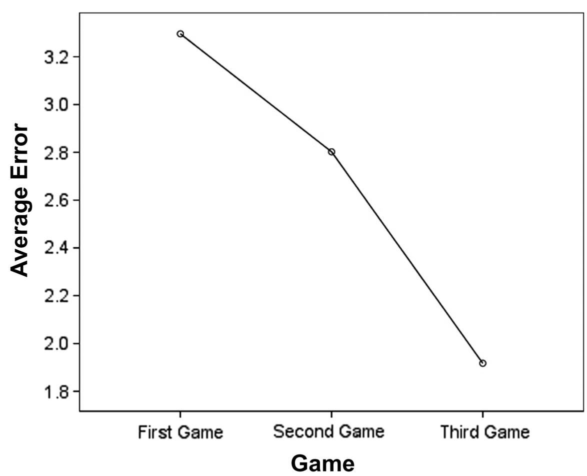

Figure 6. Average error according to sequential order.

Table 3. Observed human drumming behaviour according to order.

Order	Average error	Maximum no. of beats	Average no. of beats	Average no. of turns
1	3.30 ± 3.15	41	6.67 ± 4.22	15.88 ± 5.23
2	2.80 ± 3.36	37	5.58 ± 3.57	17.63 ± 5.84
3	1.92 ± 1.86	20	4.70 ± 2.61	19.13 ± 4.64

the preference for the third game among the participants could be explained by the lower number of errors for this game.

Behavioural data according to game type: Figure 7 shows a trend approaching statistical significance ($F(2,46) = 2.15$, $p = 0.13$) where the *gesture+* game had the highest average error, followed by the *gesture* game. The *no-gesture* game had the smallest error rate.

The maximum number of beats decreased with the increasing amount of gestures in the game (Table 4). There was a slight increase in the average number of beats with the increasing amount of gestures in the game, but this was not significant. The average number of turns tended to decrease as the amount of gestures in the game increased. This significant effect ($F(2,46) = 4.41$, $p = 0.018$) is described in Figure 8. The only interaction effect observed in this experiment between order of presentation and game type occurred for this variable ($F(2,44) = 6.020$, $p = 0.005$). This effect is described in Figure 9 and suggests that for participants who were introduced to the *gesture+* condition in the first or second game had a higher number of turns for the *no-gesture* and *gesture* game than those who encountered this game type last, while the reverse was true for the *no-gesture* condition.

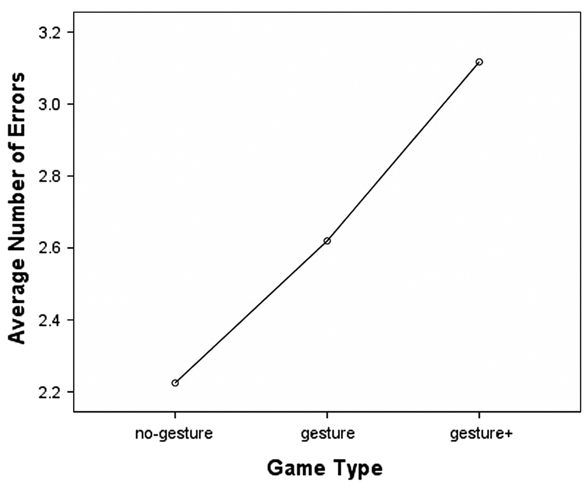

Figure 7. Average number of errors according to game type.

Table 4. Observed human drumming behaviour according to game type.

Game type	Average error	Maximum no. of beats	Average no. of beats	Average no. of turns
No-gesture	2.22 ± 2.52	41	5.24 ± 3.54	19.00 ± 5.49
Gesture	2.62 ± 3.16	37	5.60 ± 3.67	17.83 ± 4.63
Gesture+	3.12 ± 3.01	31	6.21 ± 3.89	15.58 ± 5.61

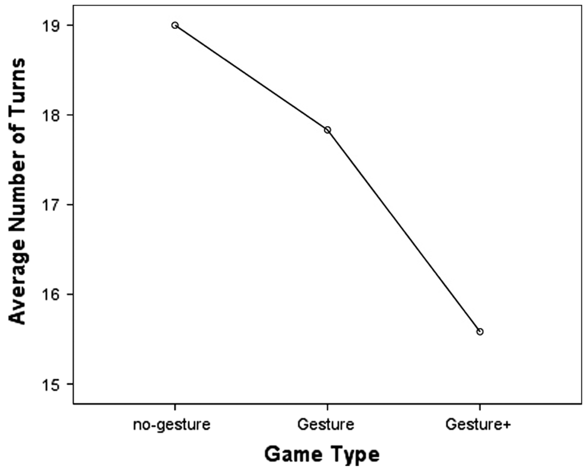

Figure 8. Average number of turns according to game type.

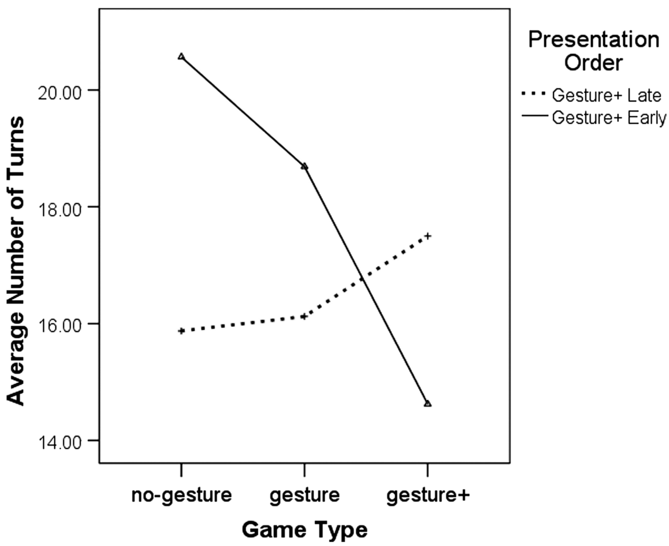

Figure 9. Interaction between game type and presentation order for number of turns.

2.4.7. *Discussion of results*

Experiment I not only investigated the possible impact of using robot gestures on drumming game (in terms of performance), but also on the participants' subsequent evaluation of the game. We expected that an intermediate level of gestures would benefit the interaction game.

Results show that the humans were indeed motivated by gestures and did, overall, enjoy the drumming experience. There did, however, seem to be a saturation level for the amount of gestures used to encourage interaction, where the amount of gestures in the gesture+ condition seemed to interfere with the participants' concentration. Drumming with no gestures, while considered efficient in terms of the drumming task, was not rated as successful in terms of social interaction. The reason for the high error rates at the start of the games is likely in part due to the participants' high expectations from the game. According to the questionnaire results, male participants appeared to view the experiment not as a game, but rather as a task to complete. Participants also may have tried to 'test' the robot's limitations during the initial stages of the trials, leading to higher error-rates, as this could involve playing rapidly in long sequences, or using different parts of the drum to create different sounds and enriching their play. They also expected that KASPAR could watch, understand and imitate them (most thought that the robot could detect them with its cameras, positioned in the eyes, and that the gestures were meaningful). As the game progressed, the understanding of the limited capabilities of the robot would increase, leading them to modify their drumming to synchronise more efficiently with the robot. This effect might have been mitigated by participant fatigue, however, as boredom was also mentioned by some participants when answering questions regarding the later games.

The data also suggest that the participants changed their style of play with the increasing level of robot gestures, playing fewer, yet longer sequences of beats.

Our sample, overall, rated the *gesture* condition as the most enjoyable, which, interestingly, had worse error rates in the evaluations of the objective performance than those without gesture. This is likely due to the *gesture* condition incorporating gestures making the interaction enjoyable to those participants who valued this aspect of the interaction, while having a lower error rate than the *gesture+* condition, and so is less adversely impacted by a task-based evaluation than this condition.

This shows that the right amount of gestures would serve to attract the attention of one portion of the participants, and make their experience enjoyable, although it did not actually help their drumming (in objective terms). This draws attention to the marked distinction between the subjective evaluations and objective performance measures.

Overall, the results from Experiment I confirmed our initial expectations (see Section 2.2), but pointed out the different effects of gesture on the dynamics of drumming performance and participants' subjective evaluation. These results helped in designing the next study (Experiment II).

3. Experiment II: emergent turn-taking

3.1. *Methodology*

As motivated earlier, one of the fundamental problems in the human–robot drumming scenario is the timing of the interaction, as timing plays a fundamental role in the regulation of human interaction. It is not always clear when the robot or human should initiate interaction in taking a turn. Therefore, in Experiment I, some predefined fixed time duration heuristics were used for synchronisation, whereby KASPAR started playing if the participant was silent for a few seconds, and would also try to motivate the participant with simple non-verbal gestures.

In Experiment II, we took a different approach and used a novel, probability-based mechanism for timing and turn-taking so that the temporal dynamics of turn-taking *emerge* from the interaction between the human and the humanoid. As explained earlier, the computational models were deliberately chosen to be simple, minimal and (algorithmically) arbitrary. Thus, these models are not meant to faithfully model turn-taking, cognition or learning in humans. Our research agenda is to study whether even such simple and arbitrary computational models will evoke different types of interaction and adaptation of people to the robot's behaviour.

We selected three different simple and minimal computational models to control the starting and stopping of the robot's regular drumming beats. This response is based on the duration time of the previous turn and on the number of beats played in the previous turn by the interaction partner. We denote the models as *Model 1, Model 2* and *Model 3*. *Model 1* uses a step function, *Model 2* a simple triangular function and *Model 3* a hyperbolic function that generates probabilities for starting or

Algorithm 1 The turn-taking algorithm

Human plays (turn # i=1)
Kaspar plays after waiting 2 seconds when human stops
FOR i=2 to *n* **DO**
  ThTime_i= $\text{KasparPlayingTime}_{i-1}$
    IF model_j ($\text{HumanPlayingTime}_i$,$\text{ThTime}_i$) = 1
      THEN KASPAR STARTS PLAYING
      ThBeat_i= # of HumanBeats_i
        IF model_j (# of KasparBeats_i,ThBeat_i) = 1
          THEN KASPAR STOPS PLAYING
END FOR (end of the game)

Figure 10. The turn-taking algorithm used in Experiment II.

stopping the robot's drumming based on these inputs from previous interaction (Figure 10). The output is bounded by maximum and minimum limits to ensure that KASPAR and the participant have time to play at least once in every turn. For every turn, the robot assesses the probability of start or stop, and takes action accordingly. For starting, the robot uses the time duration of its last bout of playing and for stopping it takes the number of beats of the human participant from the previous turn into account. The minimum number of beats KASPAR will play is one even if the resulting number of the beats recommended by any of the models is below one. The participant starts the game and KASPAR uses its turn-taking strategy when the human participant is silent for 2 s (only for the first turn). After the first turn, the turn-taking strategy is always determined by the robot's probabilistic models. Depending on the previous duration and number of beats in the interaction, according to their respective probability functions (1), (2) and (3), the return value of the three models triggers the starting or stopping in the turn-taking algorithm (Algorithm 1 in Figure 10). The probability functions for the three computational models are presented in Equations (1), (2) and (3), and visualised in Figure 11.

$$p(x) = \begin{cases} 0, & x < \text{Th} \\ 1, & x \geq \text{Th} \end{cases} \quad \text{(Step: Model 1)}, \tag{1}$$

$$p(x) = \frac{x}{\text{Th}} \quad \text{(Linear: Model 2)}, \tag{2}$$

$$p(x) = 1 - \frac{1}{x} \quad \text{(Hyperbolic: Model 3)}. \tag{3}$$

Here, x is measured in units of time for the case of starting, or, respectively, as the number of beats for stopping. Similarly, Th represents the threshold parameter of time for starting and the

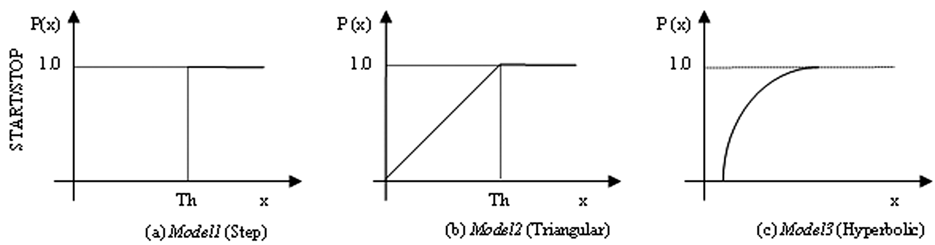

Figure 11. Computational models for START/STOP actions. For START actions, Th=ThTime, since the x-axis variable is the time (t). For STOP actions, Th=ThBeat. The x-axis variable is the number of beats (b). For START, Th is the duration of KASPAR's previous drumming bout, and for the STOP action, Th is the number of beats in the human's previous drumming bout; except that the minimum value for Th is 1.5 s (experimentally determined) for START and 1 beat for STOP actions. The only model which does not have the threshold limitations is Model 3 due to its hyperbolic nature. The y-axis gives the probability of START/STOP as a function of time/number of beats based on previous interaction.

number of beats for stopping, respectively. For each model, a decision function is called returning a 0 (‘no’) or 1 (‘yes’) is called to decide whether to change the robot’s current behaviour. That is, the function *model i* (x,Th) is called to start or stop KASPAR playing using the respective $p(x)$ function for that model as in Algorithm 1. In *Model 1,* if $p(x)$ is 1 then the model triggers starting or stopping, and this depends only on Th and the current value of x. *Models 2* and *3* have probability functions that can take values other than just 0 and 1, so a random value r in [0,1] is generated and if r is not less than the function output, then the model returns 1 (otherwise 0). Thus, in effect, in all three of these simple models, a starting or stopping action, given the current values of parameters x and Th, occurs at appropriate points with probability $p(x)$ according to the respective model, so that the model then triggers the start or stop of drumming, or otherwise no change in the behaviour occurs – see the conditionals (IF-statements) of the robot control in pseudocode of Figure 10.[8] In future, other models could also easily be assessed.

Consequently, at every turn, the robot decides when to start and stop according to the performances of both the human player and itself. Thus, the game and its dynamics are not deterministic but emerge from the moment-to-moment status of both KASPAR and the participant.

Complementary to Experiment I, we decided not to introduce any robot gestures in Experiment II but to focus our analysis on the turn-taking behaviour. Therefore in Experiment II, KASPAR did not use any gestures.

3.2. *Research questions and expectations*

In order to investigate the effect of three different generative computational models on emergent turn-taking dynamics in an imitation game, our primary research questions were as follows:

- How do different robot turn-taking strategies based on different minimal computational probabilistic models impact on the drumming performance of the human–robot pair?
- How do the different robot turn-taking strategies impact on the participants’ subjective evaluation of the drumming experience?

We expected to have ‘successful’ games in terms of turn-taking emerging from the interaction between the human and the humanoid, and that the different computational models would show different degrees of success in terms of engaging and sustaining interaction. Our ‘success’ criteria were as follows: (1) the number of turns with no or slight overlaps and gaps and (2) the number of human beats detected by the robot and the number of beats played by the robot itself that will give us hints about the quality of the games.

3.3. *Experimental conditions*

We studied three models with different parameters (Figure 11) in three different experimental conditions. We set up simulated experiments before the live experiments, to define the maximum and minimum limits and thresholds for the actual experiments with humanoid and human participants. Each model is used both for starting and stopping the robot’s play and represents an experimental condition. For *start* the time duration of the previous turn is used, and for *stop* the number of beats of the previous turn is used as a threshold. As described in the previous section, *Model 1* was a step function, where the return value of the function is ‘1’ if the input value of the function is not smaller than the threshold; thus, we expect this model to give more play time and a higher number of beats than the other models. Ideally, if the human beats long sequences, this model would reach very high values so we put a maximum time limitation (both interactants cannot play longer than 10 s per turn). Unlike *Model 1*, *Model 2* has a triangular shape which has the threshold as an upper boundary. Since we have a probabilistic approach we can have values smaller than the threshold. In fact, we expect this model to give the least play time and lowest resulting number of beats for the participants; so we foresee that the model would not be as popular as the other two models. The last condition is *Model 3*, a hyperbolic model, which cannot be limited by the thresholds. It reaches high values (close to one) very fast compared with *Model 2*. Therefore, we predict that it would result in more play time (i.e. enable the robot to play more beats than *Model 2*). Also, in our simulations, we noticed that it could enable 'coordinated games' (i.e. with a very low number of overlaps and conflicts between the human's and the robot's drumming) if we played short sequences, but since the model is not limited by thresholds, it 'reacts' to the human but does not exactly 'imitate' the human's drumming games, which we suspected that the participants might not find acceptable.

3.4. *Experiment, results and analysis*

3.4.1. *Robot*

The experiments were carried out with the humanoid robot KASPAR that was also used in Experiment I (see Section 2.4.1).

3.4.2. *Experimental setup*

The experimental setup was the same as in Experiment I (see Section 2.4.2).

3.4.3. *Software features*

The same software features were used as in Experiment I (see Section 2.4.3).

3.4.4. *Participants*

Twenty-four participants (8 female and 16 male) took part in the study. Due to logistical reasons, the trials were carried out in 2 sets (a few months apart) with 12 participants each. All participants worked in computer science or similar disciplines at the university. Only two of them had interacted with KASPAR prior to the experiment, and most were not familiar with robots in general. Six of our participants had children. (Regarding gender balance of the sample, see comment in Section 2.4.4).

3.4.5. *Interaction game setup*

We used a 1 min demonstration of the robot without any drumming game involved, where the participants were shown how to interact with KASPAR. This was followed by three games reflecting the three experimental conditions described above each lasting 3 min, without indicating to the participants anything about the differences between the conditions. The participants were simply instructed that they could play drumming games with KASPAR. As we did in Experiment I, we used all six possible different presentation orders of the games to analyse the effect of the order of the games on the humans. To account for possible fatigue, habituation or learning by the participants, in the *sequential order* section below, we analysed the games according to their order number in the sequence experienced by the participants (independent of the particular experimental condition): thus calling them the first game, second or third, disregarding their game types, for example, for one participant the first game (order 1) would be the *Model 1* game, and for another participant, *Model 1* would be the third game (order 3). After finishing the three games, the participants completed a questionnaire.

3.4.6. *Results*

3.4.6.1. Evaluation of questionnaire data. The participant evaluations were elicited in a questionnaire in the same manner as in Experiment I (see Section 2.5.1).

Most and least preferred games according to game type: See Table 5 for the number of participants which rated each game as most preferred and least preferred. There was a significant deviation from the expected counts for the most preferred game type ($\chi^2(2) = 7.76$, $p = 0.021$) as well as for the least preferred game type ($\chi^2(2) = 10.89$, $p = 0.004$). Table 5 shows that both the *Model 1* and *Model 3* games were preferred by a comparable amount of participants, while fewer participants preferred *Model 2* most.

Table 5 also shows that the highest number of the participants considered the *Model 2* game as the least preferred, while the M*odel 1* and *Model 2* games had a small number of participants which considered them the least preferred. The *Model 3* game was slightly more popular than the *Model 1* game.

Most and least preferred games according to sequential order: The number of participants which rated each game as most preferred and least preferred according to the sequential order can be seen in Table 6. The deviations from the expected count were approaching significance for the most preferred game ($\chi^2(2) = 5.25$, $p = 0.07$). Table 6 suggests that the most popular game type was the third game, while first and second games were less preferred. Table 6 also suggests that all ordinal positions of occurrence in the sequence of the games had a similar number of participants which considered them the least preferred.

As for Experiment I, in order to control for the threat against the assumptions of the counter-balancing method, mixed model ANOVAs were run using game type to investigate the possible interaction effects of presentation order and game type on both questionnaire responses and behavioural data. These were non-significant, supporting the notion of independence between presentation order and game-type overall in the sample.

Other preferences: The order of the games did not have a significant impact on the participants in terms of evaluation of the game. There were, however, significant differences according to the model used in terms of how participants evaluated the games. The participants did not rate KASPAR's drumming significantly differently across the models ($F(2,46) = 1.64$, $p = 0.20$). There was an effect approaching significance for how they rated KASPAR in terms of sociality

Table 5. Most and least preferred games according to type.

	Participants			
Game Type	Most preferred	Residual	Least preferred	Residual
Model1	9	1.7	6	−4.3
Model2	2	−6.3	17	−0.7
Model3	13	4.7	4	3.7

Table 6. Most and least preferred games according to sequential order.

	Participants			
Order	Most preferred	Residual	Least preferred	Residual
1	4	−4.0	9	0.035
2	7	−1.0	8	−0.7
3	13	5.0	9	0.035

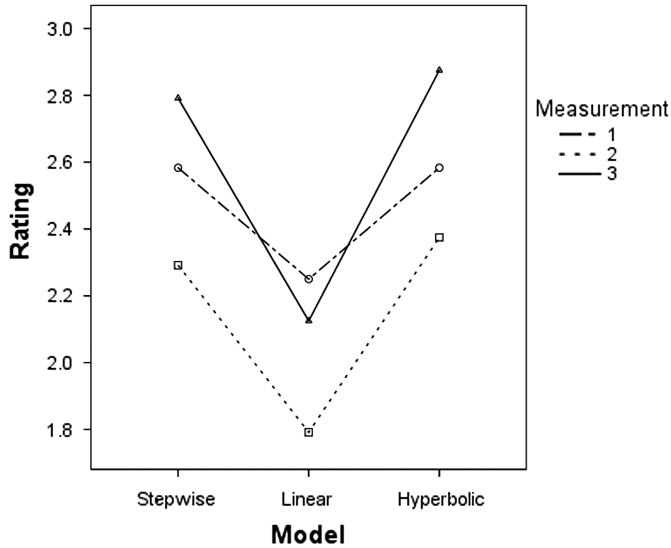

Figure 12. Ratings for the three measurements: (1) KASPAR's drumming, (2) KASPAR's perceived sociality and (3) participant enjoyment.[9]

($F(2,46) = 3.12$, $p = 0.054$), and participants significantly differentiated between the models in terms of enjoyment ($F(2,46) = 7.59$, $p = 0.001$). These effects are shown in Figure 12, which suggests that for all three, there was a tendency for the participants to rate the interactive aspects of the games lower when the linear model was used.

Sample similarities: The possibility of systematic differences between the first sample of 12 and the subsequent sample of 12 was assessed using mixed-model ANOVA. This ANOVA found no significant systematic differences between the two groups ($F(1,22) = 0.070$, $p = 0.79$). Since an identical experimental protocol was used for both groups of participants, this result supports the analysis of both samples as one larger sample. **Q5**

3.4.6.2. Evaluation of behavioural data. The behavioural data regarding the performance of the human partner during the games consisted of KASPAR's own detection of the human's drumming (denoted as 'KASPAR's view'), and video recordings of the human's drumming that were annotated and analysed manually (referred to as 'human's view'). The behavioural data includes the number of *zero turns* (where KASPAR could not register any beat performed by the human partner but played at least one beat, and passed the turn to the human), *non-zero* turns (KASPAR would register at least one drum beat of the human participant), the number of drum beats performed by human participant and KASPAR, and turn durations (referred to as 'time' in the text).

Behavioural data according to sequential order: There was no significant difference between the games according to the order (e.g. for number of turns, $F(2,22) = 0.007$, $p = 0.99$, with ANOVA). Only the human's total number of beats per game increased with the order of the games as they got used to the scenario while they played more (Table 7, KASPAR's perspective, and Table 8, human's perspective).

Behavioural data according to game type: The game types are compared in detail in Tables 9 (human's drumming) and 10 (KASPAR's drumming).

Table 7. Observed behaviour of KASPAR according to order.

Order	Average no. of beats per turn	Maximum/minimum no. of beats	Total no. of beats	Average time per turn	Maximum/minimum time per turn	Total time
1	1.7 ± 0.8	6/1	135 ± 33	1.08 ± 0.1	3/1	97.6 ± 41
2	1.72 ± 0.8	6/1	136 ± 30	1.07 ± 0.1	3/1	96 ± 40
3	1.76 ± 0.7	7/1	138 ± 26	1.07 ± 0.1	4/1	95.8 ± 41

The repeated measures ANOVA found significant differences between *Model 2* (linear model) and the other models, across a range of variables. In terms of the total number of beats there was a marked difference in the number of beats by the human registered by KASPAR ($F(2{,}46) = 58.95$, $p < 0.001$), as well as the total beats by KASPAR ($F(2{,}46) = 470.63$, $p < 0.001$), between the models used. There was no difference, however, between the models in terms of the actual number of beats *played* by the human participants ($F(2{,}46) = 0.037$, $p = 0.96$). Referring to Figure 13, we can see the relationship between detected human beats, beats *produced* by KASPAR and *the* actual *beats played by the participants across* the models.

The graph suggests that while the actual number of beats played by the humans remains more or less constant across the models, the registered number of beats decreases dramatically between the stepwise model and the other two models, while the number of beats by KASPAR increases. Thus, in the cases of linear and hyperbolic models KASPAR appeared less responsive to the playing of the participants. This result may account for the participants' higher evaluation scores for the stepwise model, compared with the linear model.

Significant differences were found between the models in terms of the ratio of turns in which KASPAR registered the beats from the human participant to the total number of turns ($F(2{,}46) = 77.18$, $p < 0.001$), see Figure 14.

Figure 15 suggests that KASPAR registered more human activity in terms of turns with both the stepwise and the hyperbolic models than with the linear model. According to Table 9, this is also clear in terms of the actual number of non-zero turns, despite the much higher number of total turns with the linear model. The difference in the actual number of turns was highly significant as well ($F(2{,}46) = 28.78$, $p < 0.001$). The above results suggest that in terms of turn-taking, KASPAR was more 'aware' (in terms of detection of beats) of the participants' behaviour in the stepwise and hyperbolic conditions than in the linear condition. The time spent drumming by the participant as registered by KASPAR may also serve to differentiate between the linear models and the two other models. There were significant differences between the three models ($F(2{,}46) = 1897.71$, $p < 0.001$), see Figure 15.

Figure 15 suggests that the amount of time in which KASPAR registered the human participant as drumming differs dramatically across the three models. The stepwise model is the most effective in this sense, followed by the hyperbolic model with the linear model being the least efficient.

These measures do suggest that some of the participants' preferences for the stepwise and hyperbolic model can be explained by objective measures of KASPAR's responsiveness to the actual drumming of the human participants. They do not, however, explain why the participants equated the stepwise and hyperbolic models in terms of enjoyment.

3.5. *Discussion of results*

Overall, the results confirm our initial expectations, namely that different computational models will lead to different human–humanoid drumming interactions (as evaluated subjectively and objectively).

Table 8. Observed drumming behaviour of human according to order.

Order	No. of turns	No. of non-zero turns	Maximum no. of beats	Total no. of beats (KASPAR's view)	Total no. of beats (human's view)	Average time per turn	Maximum/minimum time per turn	Total time
1	93 ± 44	28 ± 13	5	43 ± 23	110.2 ± 35	0.99 ± 0.6	3.11/0.01	70 ± 27.3
2	92 ± 42	29 ± 16	4	45 ± 29	113.7 ± 39	0.99 ± 0.6	2.06/0.01	69.7 ± 27.2
3	92 ± 44	31 ± 17	5	52 ± 34	115.7 ± 38	0.99 ± 0.6	3.11/0.01	68 ± 25.4

Table 9. Observed behaviour of human's drumming according to game type.

Game type	No. of turns	No. of non-zero turns	Maximum no. of beats per turn	Sum of beats (KASPAR's view)	Total no. of beats (human's view)	Average time per turn	Maximum/minimum time per turn	Total time
Model 1	66.2 ± 4	36.4 ± 15.7	5	66.7 ± 30.3	113.8 ± 33.1	1.52 ± 0.02	3.1/1.5	101 ± 5.3
Model 2	152.1 ± 3.1	21.1 ± 12.3	3	25.5 ± 14.8	114.2 ± 44	0.25 ± 0.01	0.61/0.01	37.3 ± 1.3
Model 3	58.7 ± 1.4	$30.5 + 13.8$	5	48.29 ± 23.9	111.5 ± 33.3	1.2 ± 0.01	1.8/1	70 ± 1.7

Table 10. Observed behaviour of KASPAR's drumming according to game type.

Game type	Average no. of beats per turn	Maximum/minimum no. of beats	Total no. of beats	Average time per turn	Maximum/minimum time per turn	Total time
Model 1	1.47 ± 0.28	5/1	96.5 ± 12.4	1.02 ± 0.03	3/1	67.5 ± 3.2
Model 2	1.01 ± 0.01	3/1	154 ± 3.42	1 ± 0.003	3/1	152 ± 2.8
Model 3	2.69 ± 0.1	7/2	157.9 ± 3.6	1.19 ± 0.04	4/1	69.5 ± 1.8

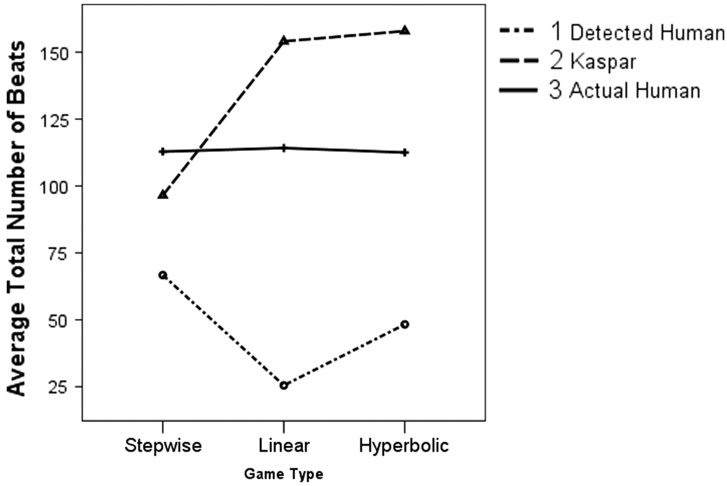

Figure 13. Total number of beats for (1) detected human beats (2) KASPAR's beats and (3) actual human beats.

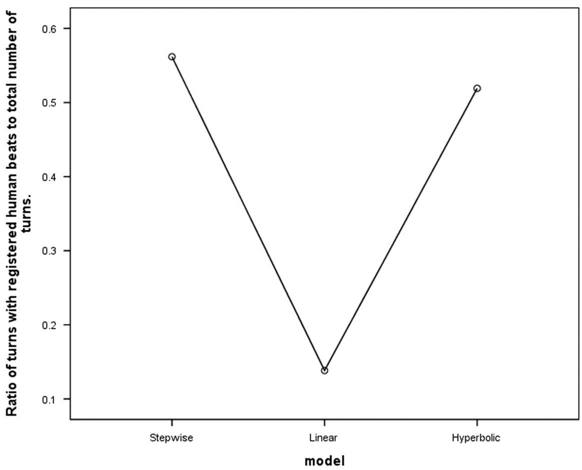

Figure 14. Ratio of turns with registered human beats to total number of turns according to the model.

As stated in the previous section, *Model 2* gave the least play time to the human and KASPAR. The impression that the participants may have got is one where KASPAR did not seem to imitate the human participants' game at all, but rather 'played on its own' (KASPAR would play at least one beat even when it did not detect a response from the human participant; Figure 16). As a consequence, KASPAR acted as a leader in the game most of the time. There were also many overlaps between KASPAR's play turns and the human participants' play turns in *Model 2*. This could be because either KASPAR or the human participants interrupted each other. More importantly, this would also cause the loss of detection of humans' beats (as described above, KASPAR would not 'listen' when it was playing). Replies to the open-ended questions in the post-game questionnaires related to this game described KASPAR's behaviour using the terms like 'annoying' or 'rude'. Thus, both the behavioural data as well as the questionnaire results describe an interaction in which the interaction's rules for turn-taking was not apparent to the

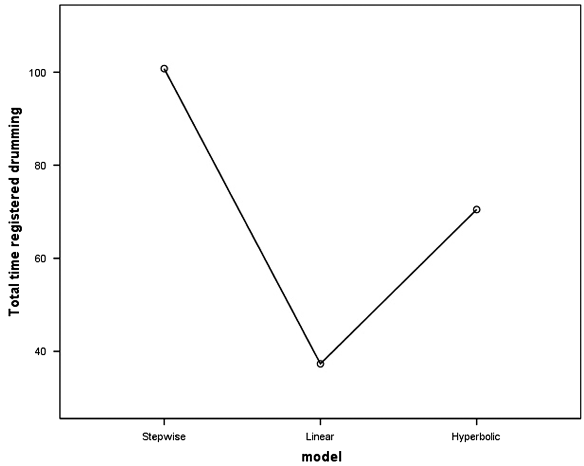

Figure 15. Total time registered for drumming according to the models.

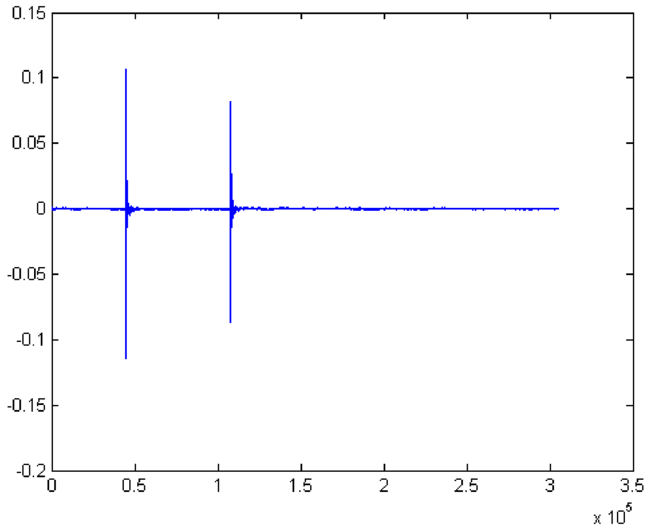

Figure 16. Representation of two beats in an example sound wave. **Q6**

human participant leading to repeated breakdowns in the social interaction, which in a human–human interaction would be described as impolite and a source of stress. Together, these measures provide an explanation as to why the participants disliked the *Model 2* game.

As stated in the previous sections, since *Model 1* uses the previous play time as a threshold, it ensures that the current play time is at least as long as the previous play time for the human player. This longer play time (compared with other games) led to both players playing longer turns which may have created the impression that the tempo of the game was slower than in the other games. This could explain the preferences for *Model 3* since the tempo of this game would be experienced as faster, having more exchanges and being perceived as more interactive. While the observed play time for the human participants was shorter than for *Model 1*, it was still long enough to allow for a coordinated game. This, coupled with the emergent nature of KASPAR's drumming in *Model 3*, led it to being viewed as more 'natural' by participants. In this game, both the human and the KASPAR played 3–4 beats in every turn (the model's probability distribution favours

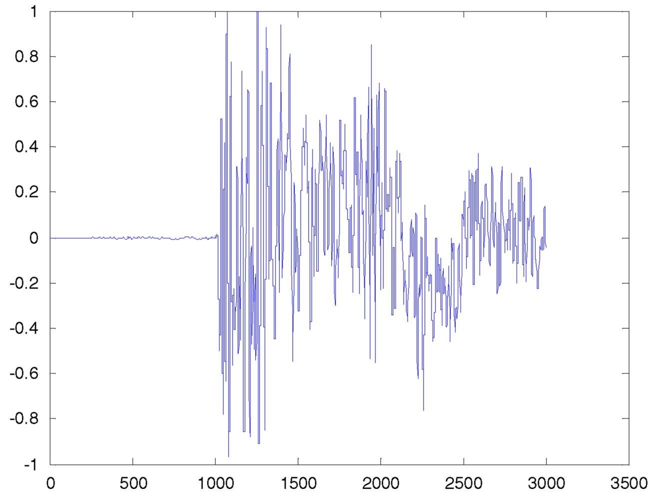

Figure 17. A single drumming bout is presented in detail. It is not a single peak but consists of many local minima/maxima.

high values), with fewer gaps in-between the participants' drumming compared with *Model 1*, and far fewer overlaps compared with *Model 2* between two turns. But *Model 3* was not bound by thresholds by nature, so it seemingly exhibited a degree of independence in regards to the human participants' performance, which some of the participants reported as being annoying. Some participants, however, did express a liking for this, though, for example, one participant described this phenomenon like 'teaching her son to play a drum'. Similarly, another participant asked if she should consider KASPAR as a professional drummer or a child while she commented on the games, since it 'looks like a child drumming rather than a professional' (Figure 17). Statements like this support the notion, suggested by the quantitative data, that the emergent turn-taking of *Model 3* was perceived to have more in common with a human–human interaction than that of the other models.

In *Model 1* the human participant was given more play time than KASPAR, but KASPAR played more beats than the human participants. However in *Model 3*, KASPAR and the participant were given almost equal durations and opportunities to play. So in the case of *Model 3*, KASPAR could act equally as a follower as well as a leader and thus had more impact on the play and played longer rhythms.

One should also note that there is a considerable amount of zero turns in all the three models. However, only in the case of *Model 2* was this amount high enough to affect the overall game. When these turns were distributed among normal turns as in *Model 1* and *Model 3*, they did not dominate the behaviour but were compensated for by the non-zero turns. But for *Model 2*, zero turns seemed to dominate the whole game and were described by the participants as a source of dislike for the model/game type.

Although there were gaps between the humans' and the robot's turns in *Model 1*, while in *Model 3* KASPAR did not seem to imitate the human participants in every turn, both models were successful in terms of emergent turn-taking. As a consequence, according to the participants' questionnaire feedback, they preferred *Model 1* and *Model 3* to *Model 2*.

As seen in the previous study, the participants actively explored the limits of KASPAR's drumming as well as the rules of the game, and adapted themselves to the games over time, which resulted in better games in terms of turn-taking and synchronisation in the later games. Thus, we observed longer sequences of playing without any overlaps or gaps between the turns. This suggests that the human participants were not passive participants in this game, but actively adapted themselves to the capabilities of the robot on their own initiative. This finding is consistent with the notion of recipient design, a concept from ethnomethodology, where we find that natural speech is always designed for its recipient (i.e. the interaction partner) and interpreted as having been so designed. Here, the speaker creates his or her turn 'with recipients in mind, and listeners are motivated to "hear" a turn that is for them and all participants closely and constantly track the trajectory of the talk to hear "their" turn' (Boden 1994, p. 71). According to conversation analysis, this turn-taking is integral to the formation of any interpersonal exchange (Boden 1994, p. 66). While in our study the robot's behaviour was controlled and based on simple computational models, we found that the participants used their recipient design skills in the interaction. The issue of recipient design will be explored further in our future research.

4. Conclusions

This article presented basic research into the regulation of interaction dynamics during social/playful HRI. We introduced an experimental setup based on human–humanoid drumming games as a suitable scenario for HRI research on non-verbal cues, synchronisation, timing and turn-taking using drumming games. Generally, the results showed that believable and enjoyable human–humanoid interaction dynamics can be created with minimal models underlying the robot's turn-taking behaviour.

Specifically, the results from this experiment suggest that there was active adaptation on the part of the participants, throughout the games. However, the efficiency of such adaptation may be countered by the participant fatigue/boredom reported in the later games, which highlights the essential role of research into how to maintain a user's interest in the interaction with a robot. One should note, however, that the results also indicate a trade-off between the subjective evaluation of the drumming experience from the perspective of the participants and the objective evaluation of the drumming performance, as well as individual differences in how the participants approached the game. The participants as a whole preferred a certain amount of robot gestures as a motivating factor in the drumming games that provided an experience of social interaction. However, the sample was divided in terms of what degrees of gestures were appropriate. The results highlight the need to ascertain to what degree the strategies used by a robot to encourage and maintain interest in such interactions, interfere with the task the interaction is centred around, as well as consider the role of individual differences in the appropriateness of these strategies. Experiment II showed that the different minimal, probabilistic models that controlled the robot's interaction dynamics led to different subjective evaluations by the participants and different dynamics in the performances of the games. The results from the questionnaires and behavioural data analysis suggest that the participants preferred the models which enable the robot and human to interact more and provide turn-taking closer to 'natural' human–human conversations, despite the differences in objective measures of drumming behaviour.

Overall, the results from our studies are consistent with the *temporal behaviour matching hypothesis* (Robins et al. 2008) which concerns the effect that the participants adapt their own interaction dynamics to the robot's. Note that our child-sized robot KASPAR, despite some human-like features such as a face, arms and few facial expressions, is still mechanical in nature (e.g. the movements are not following the biological models of movement generation, the facial expressions are minimal and not based on the models of human facial expressions, and in terms of its appearance the robot has a slightly cartoon-like appearance where we deliberately did not cover up metals and wires, e.g. protruding from the neck and wrists). But participants still adapted to the dynamics of this robot which highlights the importance of considering interaction kinesics in HRI design in general, not only in research attempting to exactly copy human-like appearance and behaviour.[10] A systematic study of the impact of robot appearance on participants' behaviour in human–humanoid drumming experiments is an interesting area of research but goes beyond the scope of the current article.

There are several noteworthy limitations of this work including methodological as well as technological limitations. Ideally, in order to generalise the results towards a wider user group the study could be repeated with participants of different age ranges, personality traits, cultural background, gender, etc. Such studies would help to explain group differences (e.g. concerning why the subjective evaluation of the participants in our study differed). Different subjective rating scales could be used. Qualitative analysis of the human–robot behavioural data (e.g. by using conversation analysis)[11] could flesh out further details of the interaction. The timing algorithms used in Experiment II could be refined in future work alongside a systematic variation of different types of robot gestures in order to find out which of these gestures have the most impact on the interaction. It may also be interesting to replicate the experiment with a different robot that had a broader spectrum of possible drumming behaviours, as this may not only enrich the interaction but also provide additional data for the performance evaluation. Last but not the least, an electronic drum could be used in order to ease the detection of the beats.

5. Future work

The HRI experiments presented in this article were based on a drumming scenario and we found that this is a very suitable task for the study of HRI and adaptive behaviour. However, our long-term research aims to go beyond a simple drumming synchronisation task and to develop richer social interaction between the robot and the human partner, which would not simply focus on synchronisation to produce the same tempo, but could provide a successful (in terms of the task) as well as enjoyable social experience to people, while allowing us to gain insight into the role of non-verbal interaction kinesics in sustaining and regulating HRI.

Based on these results, future work will investigate further issues related to interaction kinesics in general, and recipient design in particular. As mentioned above, several factors regarding robot non-verbal gestures as well as computational models underlying the robot's turn-taking behaviour seem to influence the objective performance and subjective evaluation of the interaction experience. Future work needs to investigate these further, including other factors such as the consideration of individual participants' preferences, personality profiles, as well as long-term effects.

In light of our promising results from using gestures, we foresee a system wherein KASPAR's behaviour may be motivated and rewarded by the human partner, through the partner's gestures and other expressive actions, and respond to these by playing novel acoustic rhythms and using its own repertoire of expressions and gestures to provide feedback to the human interaction partner, and, importantly, to become a 'partner' in the interaction that is not only responding but also taking the initiative proactively. If our results can be extrapolated, then such a system will be capable of motivating and sustaining interaction.

One interesting direction for future work concerns eye gaze, which plays an important role in regulating human–human interaction and communication (e.g. Kendon 1967; Farroni, Johnson, and Csibra 2004), and possibly also HRI kinesics (Mutlu, Shiwa, Kanda, Ishiguro, and Hagita 2009). While the study of gaze cues goes beyond the scope of the article, in our future work we aim to study the role of eye gaze (mutual gaze, eye gaze direction, etc.) in HRI games.

Research on interaction kinesics, as exemplified in this work, can potentially contribute to a wide range of application areas of social robots, in particular those that require long-term and repeated interaction (e.g. robots as assistive companions in the home, or robots as therapeutic or educational playmates for children). In such situations, the social acceptance of the robot, including the users' enjoyment of the interaction as well as the performance of the system in collaborative tasks, is crucial to the success of a particular application.

Acknowledgements

This work was conducted within the EU Integrated Project RobotCub ('Robotic Open-architecture Technology for Cognition, Understanding, and Behaviours'), funded by the EC through the E5 Unit (Cognition) of FP6-IST under Contract FP6-004370. We thank our colleagues in the Adaptive Systems Research Group for their help in preparing the experiments and the paper: Michael L. Walters, Joe Saunders, N. Assif Mirza, Qiming Shen, Aris Alissandrakis, Nuno Otero and Ester Ferrari.

Notes

1. In this article, we use the terms 'interactant' and 'interaction partner' (or 'partner' in short) synonymously. Thus, the term 'partner' does not imply a long-term relationship or affective bonding between human and robot.
2. KASPAR stands for kinesics and synchronisation in personal assistant robotics. The robot has been developed by our research group.
3. Note: KASPAR has previously been used successfully in studies involving children (Robins et al. 2008), including children with special needs (e.g. Robins, Dautenhahn, and Dickerson 2009) as well as adults (Kose-Bagci, Dautenhahn, Syrdal, and Nehaniv 2007; Kose-Bagci, Dautenhahn, and Nehaniv 2008). The work presented in this article is focussed on adult participants.
4. In this article, the terms 'deterministic' and 'probabilistic' turn-taking refer to the robot's control algorithm, that is, whether the robot behaves according to a deterministic or probabilistic algorithm that determines how it responds in a given moment given its sensory input. This point deserves clarification since any interaction involving humans has non-deterministic interaction dynamics as far as the *overall* human–humanoid interaction dynamics is concerned, since one cannot predict exactly how humans will behave in the interaction.
5. Preliminary results from the first 12 participants were summarised in Kose-Bagci et al. (2007).
6. Preliminary results with only an initial analysis based on 12 of the 24 participants are presented in Kose-Bagci et al. (2008).
7. KASPAR needed at least 0.3 s between beats to get its joints ready, so that, even if the human played faster, KASPAR's imitations still required minimum pauses of at least 0.3 s between the beats.
8. Note: We had also tried to *start* using beats and *stop* using time with simulated data, but the current combination resulted in more drumming time and a higher number of beats for both human and KASPAR, so this combination was preferred in the current implementation.
9. See footnote to Figure 6 above
10. For example, see android research (MacDorman and Ishiguro 2006) or other studies into the importance of robot appearance in HRI experiments (e.g. Walters et al. 2008).
11. See Robins, Dickerson, Stribling, and Dautenhahn (2004b) for an example of using conversation analysis in HRI research.
12. Similar to Kotosaka and Schaal (2000). Synchronised robot drumming by neural oscillator. *International Symposium on Adaptive Motion of Animals and Machines.*

References

Bernieri, F.J., and Rosenthal, R. (1991), 'Interpersonal Coordination: Behaviour Matching and Interactional Synchrony', in *Fundamentals of Nonverbal Behavior*, eds. R.S. Feldman and B. Rimé, Cambridge: Cambridge University Press, pp. 401–432.

Bickmore, T.W., and Cassell, J. (2005), 'Social Dialogue with Embodied Conversational Agents', in *Natural, Intelligent and Effective Interaction with Multimodal Dialogue Systems*, eds. J. van Kuppevelt, L. Dybkjaer and N. Bernsen, New York: Kluwer Academic, pp. 23–54.

Bickmore, T.W., and Picard, R.W. (2005), 'Establishing and Maintaining Long-Term Human Computer Relationships', *ACM Transactions on Computer-Human Interaction*, 12(2), 293–327.

Blow, M.P., Dautenhahn, K., Appleby, A., Nehaniv, C., and Lee, D. (2006), 'Perception of Robot Smiles and Dimensions for Human-Robot Interaction Design', in *Proceedings of the 15th IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN06)*, pp. 469–474.

Boden, D. (1994), *The Business of Talk: Organizations in Action*, Polity Press. **Q7**

Breazeal, C. (2002), *Designing Sociable Robots*, MIT Press.

Breazeal, C. (2003), 'Towards Sociable Robots', *Robotics and Autonomous Systems*, 42(3–4), 167–175, T. Fong (guest ed.).

Breazeal, C. (2004), 'Social Interaction in HRI: The Robot View', *IEEE Transactions on Systems, Man, and Cybernetics, Part C: Applications and Reviews*, 34(2), 181–186.

Condon, W.S., and Ogston, W.D. (1967), ‘A Segmentation of Behavior’, *Journal of Psychiatric Research*, 5, 221–235.

Crick, C., Munz, M., and Scassellati, B. (2006), ‘Synchronization in Social Tasks: Robotic Drumming’, *Proceedings of IEEE RO-MAN 2006*, pp. 97–102.

Dautenhahn, K. (2007a), ‘Methodology and Themes of Human-Robot Interaction: A Growing Research Field’, *International Journal of Advanced Robotic Systems*, 4(1), 103–108.

Dautenhahn, K. (2007b), ‘Socially Intelligent Robots: Dimensions of Human–Robot Interaction’, *Philosophical Transactions of the Royal Society B: Biological Sciences*, 362(1480), 679–704.

Dautenhahn, K., and Billard, A. (2002), ‘Games Children With Aautism Can Play With Robota, A Humanoid Robotic Doll’, in *Universal Access and Assistive Technology*, eds. S. Keates, P.M. Langdon, P.J. Clarkson and P. Robinson, London: Springer-Verlag, pp. 179–190.

Q8 Degallier, S., Santos, C.P., Righetti, L., and Ijsperet, A. (2006), ‘Movement Generation Using Dynamical Systems: A Humanoid Robot Performing a Drumming Task’, in *Proceedings of the IEEE-RAS International Conference on Humanoid Robots (HUMANOIDS06)*.

Farroni, T., Johnson, M.H., and Csibra, G. (2004), ‘Mechanisms of Eye Gaze Perception During Infancy’, *Journal of Cognitive Neuroscience*, 16(8), 1320–1326.

Goldin-Meadow, S., and Wagner, M.S. (2005), ‘How our hands help us learn’, *Trends in Cognitive Science*, 9(5).

Goodrich, M.A., and Schultz, A.C. (2007), ‘Human-Robot Interaction: A Survey’, *Foundations and Trends in Human-Computer Interaction*, 1(3), 203–275.

Hall, E.T. (1983), *The Dance of Life: The Other Dimension of Time*, Doubleday: Anchor Press.

Hendriks-Jansen, H. (1996), *Catching Ourselves in the Act Situated Activity, Interactive Emergence, Evolution, and Human Thought*, Cambridge, MA: MIT Press.

Iizuka, H., and Ikegami, T. (2004), ‘Adaptability and Diversity in Simulated Turn-Taking Behavior’, *Artificial Life*, 10(4), 361–378.

Ito, M., and Tani, J. (2004), ‘Joint Attention Between a Humanoid Robot and Users in Imitation Game’, in *Proceedings of the 3rd International Conference on Development and Learning (ICDL’04)*, La Jolla, USA.

Kendon, A. (1967), ‘Some Functions of Gaze Direction in Social Interaction’, *Acta Psychologica*, 26, 22–63.

Kendon, A. (1970), ‘Movement Coordination in Social Interaction: Some Examples Described’, *Acta Psychologica*, 32, 100–125.

Kose, H., and Akin, H.L. (2001), ‘Object Recognition in Robot Football Using a One-Dimensional Image’, in *The Tenth Turkish Symposium on Artificial Intelligence and Neural Networks (TAINN 2001)*, pp. 291–300.

Kose-Bagci, H., Dautenhahn, K., Syrdal, D.S., and Nehaniv, C.L. (2007), ‘Drum-Mate: A Human-Humanoid Drumming Experience’, in *IEEE-RAS Humanoids2007*.

Kose-Bagci, H., Dautenhahn, K., and Nehaniv, C.L. (2008), ‘Emergent Turn-Taking Dynamics in Drumming Games with a Humanoid Robot’, in *IEEE RO-MAN 2008*.

Kotosaka, S., and Schaal, S. (2001), ‘Synchronized Robot Drumming by Neural Oscillator’, *Journal of the Robotics Society of Japan*, 19(1), 116–123.

Kozima, H., Nakagawa, C., Yasuda, Y., and Kosugi, D. (2004), ‘A Toy-Like Robot in the Playroom for Children with Developmental Disorder’, in *Proceedings of the International IEEE Conference on Development and Learning (ICDL-04)*, San Diego, USA.

MacDorman, K.F., and Ishiguro, H. (2006), ‘The Uncanny Advantage of Using Androids in Social and Cognitive Science Research’, *Interaction Studies*, 7(3), 297–337.

Metta, G., Fitzpatrick, P., and Natale, L. (2006), ‘Yarp: Yet Another Robot Platform’ (special issue on Software Development and Integration in Robotics), *International Journal on Advanced Robotics Systems*, 3(1), 43–48.

Michalowski, M.P., Sabanovic, S., and Michel, P. (2006), ‘Roillo: Creating a Social Robot for Playrooms’, in *Proceedings of IEEE RO-MAN 2006*, pp. 587–592.

Miller, S.H. (1984), *Experimental Design and Statistics*, London: Routledge.

Miyake, Y. (2003), ‘Co-creation in Man-Machine Interaction’, in *Proceedings of the 12th IEEE International Workshop on Robot & Human Interactive Communication (RO-MAN03)*, IEEE Press, pp. 321–324.

Mutlu, B., Shiwa, T., Kanda, T., Ishiguro, H., and Hagita, N. (2009), ‘Footing in Human-Robot Conversations: How Robots Might Shape Participant Roles Using Gaze Cues’, in *Proceedings of the 4th International Conference on Human-Robot Interaction (HRI09)*, March 2009, San Diego, CA.

Nadel, J., Guerini, C., Peze, A., and Rivet, C. (1999), ‘The Evolving Nature of Imitation as a Format of Communication’, in *Imitation in Infancy*, eds. J. Nadel and G. Butterworth, Cambridge University Press, pp. 209–234.

Nass, C., and Lee, K.M. (2000), ‘Does Computer-Generated Speech Manifest Personality?’ in *Proceedings CHI2000*, pp. 329–336.

Pfeifer, R., and Scheier, C. (1999), *Understanding Intelligence*, MIT Press.

Pickering, M.J., and Garrod, S. (2004), ‘Toward a Mechanistic Psychology of Dialogue’, *Behavioral and Brain Sciences*, 27, 169–225.

Reeves, B., and Nass, C. (1997), *The Media Equation: How People Treat Computers, Televisions, and New Media Like Real People and Places*, New York: Cambridge University Press.

Robins, B., Dautenhahn, K., te Boekhorst, R., and Billard, A. (2004a), ‘Effects of Repeated Exposure of a Humanoid Robot on Children with Autism’, in *Designing a More Inclusive World*, eds. S. Keates, J. Clarkson, P. Langdon and P. Robinson, London: Springer Verlag, pp. 225–236.

Robins, B., Dickerson, P., Stribling, P., and Dautenhahn, K. (2004b), ‘Robot-Mediated Joint Attention in Children With Autism: A Case Study in Robot-Human Interaction’, *Interaction Studies*, 5(2), 161–198.

Robins, B., Dautenhahn, K., Nehaniv, C.L., Mirza, N.A., Francois, D., and Olsson, L. (2005), 'Sustaining Interaction Dynamics and Engagement in Dyadic Child-Robot Interaction Kinesics: Lessons Learnt from an Exploratory Study', in *Proceedings of IEEE RO-MAN'05*, IEEE Press, pp. 716–722.

Robins, B., Dautenhahn, K., te Boekhorst, R., and Nehaniv, C.L. (2008), 'Behaviour Delay and Robot Expressiveness in Child-Robot Interactions: A User Study on Interaction Kinesics', in *Proceedings of the ACM/IEEE 3rd International Conference on Human-Robot Interaction* (*HRI '08*), ACM, New York, pp. 17–24.

Robins, B., Dautenhahn, K., and Dickerson, P. (2009), 'From Isolation to Communication: A Case Study Evaluation of Robot Assisted Play for Children with Autism with a Minimally Expressive Humanoid Robot', in *Proceedings of the Second International Conferences on Advances in Computer-Human Interactions, ACHI 09*, 1–7 February 2009, Cancun, Mexico, IEEE Computer Society Press, pp. 205–211.

Sacks, H., Schegloff, E.A., and Jefferson, G. (1974), 'A Simplest Systematics for the Organization of Turn-Taking for Conversation', *Language*, 50, 696–735.

Schmidt, R.C., Richardson, M.J., Arsenault, C., and Galantucci, B. (2007), 'Visual Tracking and Entrainment to an Environmental Rhythm', *Journal of Experimental Psychology*, 33(4), 860–870.

Shiwa, T., Kanda, T., Imai, M., Ishiguro, H., and Hagita, N. (2008), 'How Quickly Should Communication Robots Respond?', in *ACM/IEEE 3rd Annual Conference on Human-Robot Interaction (HRI2008)*, pp.153–160.

Suzuki, N., and Katagiri, Y. (2007), 'Prosodic Alignment in Human-Computer Interaction', *Connection Science*, 19(2), 131–141.

Steels, L., and Brooks, R.A. (eds.) (1995), *The Artificial Life Route to Artificial Intelligence: Building Embodied Situated Agents*, Hillsdale, NJ: Lawrence Erlbaum Associates, Inc.

Syrdal, D.S., Koay, K.L, Walters, M.L., and Dautenhahn, K. (2007), 'A Personalised Robot Companion? The Role of Individual Differences on Spatial Preferences in HRI Scenarios', in *IEEE International Symposium on Robot and Human Interactive Communication (Ro-man)*, Jeju Island, Korea, pp. 1143–1148.

Tapus, A., and Matarić, M.J. (2006), 'Towards Socially Assistive Robotics', *International Journal of the Robotics Society of Japan*, 24(5), 576–578.

Tapus, A., Tapus, C., and Mataric, M.J. (2008), 'User-Robot Personality Matching and Robot Behavior Adaptation for Post-Stroke Rehabilitation Therapy' (special issue on Multidisciplinary Collaboration for Socially Assistive Robotics), *Intelligent Service Robotics Journal*, 1(2), 169–183.

Thórisson, K.R. (2002), 'Natural Turn-Taking Needs No Manual: Computational Theory and Model, from Perception to Action', in *Multimodality in Language and Speech Systems*, eds. B. Granström, D. House and I. Karlsson, Dordrecht, The Netherlands: Kluwer Academic Publishers, pp. 173–207.

Trevarthen, C. (1999), 'Musicality and the Intrinsic Motive Pulse: Evidence from Human Psychobiology and Infant Communication' (special issue), *Musicae Scientae*, 155–215. **Q9**

Walters, M.L., Syrdal, D.S., Dautenhahn, K., te Boekhorst, R., and Koay, K.L. (2008), 'Avoiding the Uncanny Valley: Robot Appearance, Personality and Consistency of Behavior in an Attention-Seeking Home Scenario for a Robot Companion', *Autonomous Robots*, 24(2), 159–178.

Watanabe, T. (2004), 'E-cosmic: Embodied Communication System for Mind Connection', in *Proceedings of IEEE RO-MAN'04*, Kurashiki, Japan.

Weinberg, G., and Driscoll, S. (2006), 'Robot-Human Interaction With an Anthropomorphic Percussionist', in *Proceedings of International ACM Computer Human Interaction Conference (CHI 2006)*, Montreal, Canada, pp. 1229–1232.

Weinberg, G., Driscoll, S., and Parry, M. (2005), 'Musical Interactions With a Perceptual Robotic Percussionist', in *Proceedings of IEEE International Workshop on Robot and Human Interactive Communication (RO-MAN 2005)*, Nashville, TN, pp. 456–461.

Wrede, B., Buschkaemper, S., and Li, S. 'Do You Like This Robot? The Role of Robot Behavior, Robot Personality and User Personality', in *ACM/IEEE Human-Robot Interaction Conference – HRI2006*, Salt Lake City, Utah, USA. **Q10**

Yamamoto, M., and Watanabe, T. (2003), 'Time Lag Effects of Utterance to Communicative Actions on Robot-Human Greeting Interaction', in *Proceedings of IEEE RO-MAN'03*, pp. 217–222.

Yamaoka, F., Kanda, T., Ishiguro, H., and Hagita, N. (2007), 'How Contingent Should a Lifelike Robot Be? The Relationship Between Contingency and Complexity', *Connection Science,* 19(2), 143–162.

Yamaoka, F., Kanda, T., Ishiguro, H., and Hagita, N. (2008), 'Developing a Model of Robot Behavior to Identify and Appropriately Respond to Implicit Attention-Shifting', in *Proceedings of ACM/IEEE HRI 2009 Conference*, pp. 133–140.

Yoshikawa, Y., Shinozawa, K., Ishiguro, H., Hagita, N., and Miyamoto, T. (2006), 'Responsive Robot Gaze to Interaction Partner', in *Proceedings of Robotics: Science and Systems*.

Appendix 1. Audio analysis

The acoustic sound waves recorded by the sound grabber module are converted to digital music samples, which allows the use of mathematical computations and sample-based techniques. To detect the patterns of a sound wave, a filter-based method is used, based on the work of Kose and Akin (2001) originally used to detect visual patterns. This method which is called Audio Analyser was used in the drumming experiments with KASPAR as well as a different humanoid robot (iCub) in real time. Also, in work not reported in this article, it was integrated to Webots software (Cyberbotics) to be used in a simulated drummer modelled after the iCub robot. The real power of the method comes from its being computationally efficient, simple, fast and real time. The drumming experiments are real time, and to have games which appear 'natural' with short durations between turns, we need to identify the bouts of drumming as soon as they are produced. Therefore, it is not possible to record them first and perform off-line analysis, or use efficient but complex methods in terms of computational resources and time. Also, although the human participants are expected to use either the end of a pencil or one hand to hit the toy drum, many different strategies were observed (and people were not discouraged to use them): they were observed to use the tambourine-style bells around the drum, use both hands or sometimes use a pen or a stick to hit the drum. Therefore, it is not trivial to train the system with 'normal' drumming bouts. Also, the high inner noise of the humanoid, besides the high noise around the drumming area (due to people present in the room), makes the environment very challenging and require us to set up high noise filters. The noise filters should be high enough to filter out the inner and outer noise, but low enough to pick up as many drumming bouts as possible. Since we use participants from both genders and all age groups, we could observe very frequent or very light bouts of drumming which are even harder to analyse. In the current implementation, we only use audio feedback to detect the drumming bouts, but in future work, we plan to use visual feedback also. However, as we mentioned earlier, the participants were allowed to use various different ways to produce sound during their drumming games; so even the addition of visual feedback would not bring optimal success in bout detection.

To detect the patterns inside a sound wave, a filter-based method is used.[12] In this method, a four-item mask is applied to every sample in the sound wave, and a filter is constructed. The peaks in this filter show the edges in the sound wave. A mask of $[-1\ -1\ 1\ 1]$ is used to detect rising edges, and another mask of $[1\ 1\ -1\ -1]$ is used to detect falling edges. Any part of the sound wave between a rising and a falling edge is a region which represents the beat in the sound wave. This is because a beat is represented by a set of points and not a single point. Once the regions are detected, a threshold is applied on the average value of the points in the region, to detect the 'real beats' and discard noise. This method is computationally simple but fast and efficient.

Drumming with a Humanoid Robot: Lessons Learnt from Designing and Analysing Human-Robot Interaction Studies

Hatice Kose-Bagci, Kerstin Dautenhahn, and Chrystopher L. Nehaniv

Adaptive Systems Research Group
University of Hertfordshire, School of Computer Science
Hatfield, Herts, AL 10 9AB, U.K.
{H.Kose-Bagci,K.Dautenhahn,C.L.Nehaniv}@herts.ac.uk

Abstract

We summarize methodological and experimental design issues related to three human-robot interaction studies investigating a drumming experience with Kaspar, a humanoid child-sized robot, and (in total 116) human participants. Our aim[1] is not to have Kaspar just replicate the human's drumming but to engage in a 'social manner' in a call and response turn-taking interaction. This requires the set up of enjoyable as well as (as much as possible) controlled experiments. Two Human-Robot Interaction (HRI) experiments with adult participants and one experiment with primary school children were carried out to investigate different aspects of such interactions. We briefly summarize issues concerning experimental methodology and design, as well as ethical, legal, safety issues in addition to many 'practical' challenges of setting up and conducting HRI experiments with an autonomous humanoid robot.

Introduction

We present methodological and experimental design issues related to three exploratory studies investigating a drumming experience (*drum-mate*) with Kaspar [Blow, et al., 2006] and human participants. This research is part of a project in developmental robotics with a particular emphasis of our work on gesture communication. The primary goal of this work is to achieve *(non-verbal) gesture communication* between child-like humanoid robots and human beings, whereby drumming served as a test bed to study key aspects such as turn-taking and non-verbal gestures.

In the first study turn-taking is deterministic and head gestures of the robot accompany its drumming to assess the impact of non-verbal gestures on the interaction [Kose-Bagci, et al., 2007]. The second study focuses on emergent turn-taking dynamics; here our aim is to have turn-taking and role switching which is not deterministic but emerging from the social interaction between the human and the humanoid [Kose-Bagci, et al., 2008]. Each of these two experiments were carried by 24 adult participants (in total 48 adult participants were involved). The third study with 68 primary school children focuses on the effect of embodiment and gestures on the subjective and objective evaluations of the human participants (details of the study and results will be published in a future publication [Kose-Bagci, et al. in preparation]). In all three studies (whose detailed results are reported elsewhere), participants did not have any prior experience with robots. All the experiments were carried out in real-time, and the humanoid robot was operating completely autonomously.

The remainder of this paper is organized as follows, the next section overviews the *drum-mate* studies, their methodology and the research questions motivating them. The experiment design section describes the experimental setups, provides brief information about the humanoid robot Kaspar, and the game setup, followed by a section on data collection. Legal and safety issues, as well as ethical, experimental and other methodological issues are discussed in the following sections. The last section includes a brief conclusion on the experiments, lessons learnt, and presents ideas for future work.

Drum-mate

Methodology

Drum-mate is an interactive drumming game played by a human participant and an autonomous humanoid robot.

[1] Acknowledgements: This work was conducted within the EU Integrated Project RobotCub ("Robotic Open-architecture Technology for Cognition, Understanding, and Behaviours"), funded by the EC through the E5 Unit (Cognition) of FP6-IST under Contract FP6-004370.

The game was enriched by non-verbal gestures, or different computational turn-taking models according to the research interests of the different studies. The human participant starts the game by playing a simple rhythm on his/her toy drum. When the robot 'understands' that the human has finished playing, the robot takes its turn, trying to imitate the human's drumming. Then the human takes his/her turn at drumming, and so on. This continues for a fixed amount of time, e.g. 3 minutes in the experiments involving adult participants. The robot uses audio feedback to regulate turn-taking and imitation.

Research Questions and Expectations

In the first experiment, we studied the effect of the robot's social gestures in a game of imitation and (deterministic) turn-taking. We expected that participants would be more engaged (in terms of drumming performance) and evaluate the interactions (questionnaires) more positively in the experimental condition when Kaspar used head gestures while imitating the human's drumming than when no such gestures were used [Kose-Bagci, et al., 2007].

The second experimental study investigated the effect of three different probabilistic computational models on the emergent turn-taking dynamics in a drumming game. The game was a modified version of the *drum-mate* game where Kaspar used no gestures but only drumming, and the game mainly focused not on imitation but on turn-taking dynamics emerging from the social interaction between the humanoid and the human participant. Here we expected the different probabilistic controllers to impact the interaction experience significantly. As in the first experiment, objective measures of drumming performance, as well as the subjective evaluations by the participants were analysed [Kose-Bagci, et al., 2008].

The third set of experiments mainly focused on the effect of different embodiments and non-verbal gesture conditions on the interaction between children and the humanoid robot. Like the first two experimental studies with adults, in this work, we also analyse the results in terms of performances of the robot and the human participants and subjective evaluations (questionnaires). Our research interests mainly focused on the differences between conditions where children play in real-time interaction with either the physical robot, the projection of the remotely located robot, or with the 'disembodied' robot (only the sound of the hidden robot is available to the children). Also we expected that these differences would increase in the presence of additional robot gesturing.

Experimental Design

Kaspar

The experiments were carried out with the child-like humanoid robot Kaspar which was designed and built by the members of the Adaptive Systems Research Group at the University of Hertfordshire to study human-robot interactions with a minimal set of expressive robot features. Kaspar has 6 degrees of freedom in the head and neck, 2 in the eyes that are fitted with video cameras, a mouth capable of opening and smiling, and 4 in the each arm. The face is a silicon-rubber mask, which is supported by an aluminium frame [Blow, et al., 2006]. It has immobile legs and fixed feet and hands.

Experimental Setup

The first two experiments (with adult participants) were carried out in a separate room isolated from other people and noise which could affect the drumming experiment. Kaspar was seated on a table with the drum on its lap (Figure 1). The human partner was seated in front of the robot using another drum that was placed on the table. The human participants were to use a pencil to hit the drum. [Although we suggested to the participants to use one pencil and hit the top of the drum, sometimes they used two pencils, or they used their bare hands (single hand or with both hands) and several times they used the tambourine-style bells around the drum's sides.]

The third experiment (with children) was carried out in two almost identical cubicles isolated from the rest of a room with high barriers (Figure 2). In the rest of the room other robotic activities took place at the same time with other children. In one cubicle we had Kaspar seated on a table with the drum on its lap, similar to the first two experiments. In the second cubicle instead of the table and the robot Kaspar, we projected a real time image of Kaspar on a whiteboard, in order to study the effect of the robot's embodiment (physically present versus remotely located). For the third condition, where only the sound of the robot was heard (disembodied robot), we use the same setup but turned off the projector. (See Figure 3). All other aspects of the experimental setup were kept the same.

Fig. 1 A screen shot from the experiments

(a) Physically embodied robot, Kaspar and child participant

(b) Child participant is playing the game in the *projection condition* where the robot's live performance is captured by a camera and projected to a wall in front of the child.

(c) Child participant playing the game with the 'disembodied' robot in the *sound only* condition, where the robot is hidden but audible and only its drumming sound is heard.

Fig. 2 Screenshots from the experiment with children where they played (a) with the physical robot, (b) watching a projected image of the robot, and (c) hearing only the sound of the robot

Interaction Game Setup

For the first two experiments with adults, before every experiment, for each participant, we used a one minute demo of the robot where participants were shown how to interact with Kaspar. Here the participant played the game following a brief introduction of the robot and the game from the experimenter. They learned the rules of the game and got used to the robot without being video recorded. This was followed by three games reflecting the three experimental conditions [Kose-Bagci, et al., 2007; Kose-Bagci, et al., 2008] each lasting three minutes, without indicating to the participants anything about the differences between the conditions.

For the third experiment, we again had three games with different conditions, each of which took two minutes. We had a 30 seconds demo of the first condition in that session which was carried out by one of the experimenters with necessary explanations which was same for all the sessions. So if the first condition is the one with the projected image of Kaspar, then the children see exactly the demo of this condition, not the demo with the physical Kaspar itself. Unlike the first two experiments, the demo is given to a group of children. After the demo the participants will play individually with the robot in the actual experiment.

For each experiment, we used all possible different presentation orders of the games, to analyze the effect of the order of the games on the humans. This is important for avoiding possible fatigue or learning by the participants.

Compared to the experiments with adults, in the experiments with children, simpler gestures were used, and the game duration was decreased to two minutes from three minutes. Also the time between turns was decreased to adapt the game better for the children (i.e. to make the game faster and easy to understand).

Data collection

In the experiments, data were collected to analyse how the human participants evaluate different games and both the robot's and the human participants' performances during these games. We had three main sources of data in our experiments: questionnaires, the drumming data recorded by the robot itself, and the video recordings of the trials including the human partners' drumming behaviour which were then annotated and then quantitatively analysed.

Questionnaires, Consent, Ethics

Before starting the trials, each adult participant was given a questionnaire and a consent and demographics form involving a short description of the experiment and related work. As described in detail in the following subsection, video recordings are important data sources for our studies, so in the consent forms we ask the participants' permission to record their performances by video cameras during the experiments and to use these recordings to produce photos and movies for scientific presentations. We also used these video recordings as data sources to analyse the performances of the participants and the robot. Therefore, participants were also given the option to consent to the video recording and analysis, but not to the use of videos for scientific presentations. If participants do not consent at all to be video recorded during the experiment then their data have to be excluded from the trial.

Unlike some other HRI experiments, our experiments were totally volunteer based, our participants were not paid for their attendance of the trials. This made it difficult to find a large number of participants, with different features e.g. female or left handed participants. Also in the consent forms we inform them that participation in the study is voluntary and that they can leave the experiment at any point during the experiment without being questioned.

Moreover, an ethics approval form including very detailed information about the experiments regarding safety, data collection etc. had to be submitted to the faculty ethics committee of the university. Approval had to be granted before the recruitment of participants and the actual experiments could start. In the case of child participants, in addition a parental consent form was sent to each parent of the children involving detailed information about the experiment and the presence of robots in the experiment, including the possibility of recording the sessions with video cameras and using these recordings later in scientific presentations. According to the result of these forms, some sessions were not recorded by video, or the recorded video was just used for data analysis but not in scientific presentations.

In the first two experiments with adults, the questionnaire which was given before the experiments included general questions about the adult participants, e.g. name, age, nationality, their profession, and if they are parents/careers of children (to understand if they are used to playing with children/children's games).

The children were asked different questions e.g. regarding their tendency to play video games. We were very careful about not asking questions regarding their nationality or ethnic origin, which might be interesting in scientific terms (for cross-cultural comparisons) but may offend or cause discomfort to them. Also we tried to put the questions as simply and understandably as possible, and used small pilot groups to test the usability and understandability of the questionnaires before the real experiments.
After each game each participant completed one page of the questionnaire to express her/his opinion about the game s/he had just played (evaluation of the game and the robot's behaviour in that particular game i.e. sociable/unsociable, or enjoyable/not enjoyable, quality of the interaction for child participants, and the evaluation of the game, drumming of the robot and the social interaction with the robot in the case of adult participants).

Once the participants had completed the items related to the last (3rd) game, they completed the last session of the questionnaire where they could judge the overall experience by deciding which of the three games they liked the best and which they liked the least and the reasons behind that decision.

In adults they also had to judge whether there were any differences between the three games and state these differences.

We tried to keep the number of questions very brief and used simple and direct phrases in the questionnaire.

All the questions about the evaluations were scale based, and the participants were encouraged to write down and express their detailed feelings and suggestions after each evaluation. All of our questionnaires were designed with the help of our psychologist team members for the benefit of the participants.

Video

The experiments were recorded by at least two different cameras positioned at different locations of the experimental area (one facing the human participant and another facing the robot), at each single game. Usually a second experimenter was present in the experimental area to control the cameras, start/stop them during the trials, and (in order not to waste video tape) not to record when the participants played demo games, or worked on the questionnaires between the trials. We used cameras with tapes and fixed them on tripods in several locations of the experiment room where we can view the experiment but do not interfere with the robot or the participant. We even used small tripods to fix the cameras on top of book shelves or doors to get the best viewing angle (to see both faces of the robot and human and the drums).

The video recordings were then analysed manually to detect the performance of the human participant's behavioral data (e.g. the number of drum bouts played by the human, and number of turns taken by the human at each game). Also the video recordings gave clues about the humans' behaviours at certain situation, how they reacted physically and emotionally in different conditions. These observations and the data gathered were compared with the questionnaire results and the data recorded by the robot itself.

Robot

The behavioural data taken from the robot itself includes clues about the robot's and the human participant's performance. Kaspar records its performance (e.g. the number of drum bouts played by Kaspar, and the number of turns taken by the Kaspar at each game). Kaspar also records the human's performance as well (e.g. the duration between two drum bouts) to imitate their performance within its physical limitations. There are two main sources of behavioral data: the audio feedback taken by microphones and analysed in real-time by the robot to extract features of the drumming (i.e. number of drum bouts played, timings and the durations between drum bouts), and the robot's own actions (i.e. name of the action taken at that turn, joint values that are activated during the actions, and timings of the actions).

Note that Kaspar does not have a memory unit on board so this data is recorded in real time on a laptop, not on Kaspar itself .

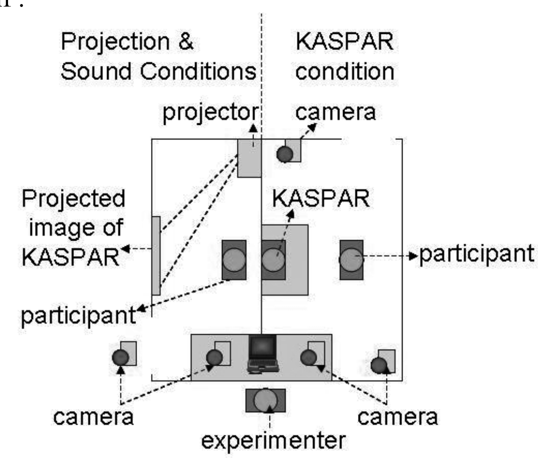

Fig. 3 The experimental layout of the experiment with children playing drumming games with Kaspar in different conditions that varied in terms of the robot's embodiment.

Note, the robot's gestures were kept very simple (e.g. simple head moves, nodding and blinking in the experiments with adults, and additionally a smile to show 'happiness' and a neutral smile to show 'sadness' when the robot could not get any feedback from the human, and waving a hand to say the game is over, for the children participants). More 'expressive' gestures might have distracted participants and could have interfered with our research agenda.

Legal and safety issues

In our experiments we worked with only one humanoid**,** Kaspar, which was risky in case the robot broke down. Its power supplies were low voltage (6V and 12V) rechargeable lead-acid gel batteries for safety. The batteries should be charged fully before each trial, especially if the robot was very active during the experiments. Therefore we limited the use of the robot to a few hours a day. We always had access to a researcher responsible from the maintenance of the robot. The robot never interacted physically with the human, e.g. by touching, and it was stable on the table and did not move its body during the trials, except its head and arms. Kaspar was placed on a table and we always kept a safety distance (at least 30 cm) between the robot and the participants.

All of our researchers who worked with children had CRB clearance (CRB- Criminal Records Bureau, UK). During the event involving children visiting our University (with an opportunity of participants to experience various demonstrations of interactive software and robots, see [FearNot! event, 2008]), their teacher, and a psychologist from our team accompanied the children. The questionnaires were prepared and given to the children by the psychologist. Neither adults nor children were left alone with the robot. The children were not allowed to have sweets during the whole event, but had lunch which was prepared according to their dietary requirements.

Ethical Issues

Although our experiments targeted specific research questions they were also designed to be 'fun' for the participants who volunteered to take part in the study. This was particularly important for the trials involving children who visited our university as part of a school excursion that was meant to have an educational but also enjoyable nature. Thus, we had to use a limited number of experimental conditions in order not the make the event boring for them. For the same reason the duration of the experiments had to be kept short, and overall the experiments had to be designed in a 'pleasant' manner (participants should feel relaxed and comfortable during the experiment). Our intention was to create a playful and engaging setup for the study of human-robot interaction, so many aspects of the work could not be as rigorously controlled as may have been desirable from a purely experimental design point of view. These ethical issues are important to consider in particular for children and other vulnerable people, since researchers do not want to "waste the time" of participants who volunteer in our studies and thus contribute to our research.

Experimental and other methodological issues

Before the experiments are set up, the experimental area and all equipment involved needs to be checked carefully, e.g. check that the video cameras and robots are working properly etc. Also, once the experimental design and setup have been decided, running a simulation of the experiment (whereby experimenters may take the role of participants) is important to see if the time restrictions are satisfied. This is a must especially when running experiments with lots of participants on the same day. Especially in the experiment with children, we had to consider the timing for each child, including their preparation (entering the experimental area, sitting down etc.), the duration of the actual interaction experiment and the time required to complete the questionnaires. In addition, the experimenters need to prepare and practice explanations and answers to possible questions from the children about the work, including questions about the robot's functioning etc. Note, this requires to a) provide as much detail as necessary in order to satisfy the curiosity of the children, and b) not to disclose too many details that may confuse or overwhelm the children, or introduce strong biases influencing the outcome of the experiments.

We also decided to work with the children in teams, not alone, which would help them to get used to the robot and the game. But working with teams of children lead to other issues, like how to identify the children individually (and match the ID codes that they were given to the results of the data collected during the experiment etc.).

In addition to within-subject comparisons (e.g. for testing different game conditions and different gestures of the robot) we also carried out between-subjects tests in order to study the impact of individual features of the participants which may affect the results, e.g. gender.

Our experiments were 'controlled' (in the sense that each participant in a particular study was exposed to a specific and clearly defined experimental set up whereby a particular experimental procedure was followed involving different experimental conditions) but they were not fully restricted laboratory experiments because we intended to create a 'playful atmosphere' in order to facilitate natural interaction that is emerging between the robot and the human and any adaptation of the human to the robot and the game. Also, all our participants were volunteers so we needed to achieve an enjoyable task as well as a controlled experiment. Last but not least, the group of children we worked with was an opportunity sample, and while we could control for age (due to the fact that the children arrived as part of a school class) we could not control e.g. for gender.

Audio analysis was a very vital part of our game, and the robot's "hearing" was effected by both external and internal noise (coming from the robot's motors) so we had to use some noise filters or cover our microphones especially in the experiment with the children where we had 10-20 children shouting and talking in the same room (the experimental area was only separated by screens). Note, in the experiments with the children we deliberately decided to carry out the experiments in the same room were the other robotic group activities took place. Alternatively, we had considered isolating individual children or small groups of children and leading them to a different nearby experimental room. However, we decided against such a more 'controlled' approach since it would have interfered with the enjoyable nature of the event (as part of a school excursion in order to learn, playfully, about robots and virtual characters). It also ethically did not seem to be justified to remove individual children from their peer group in this group-oriented event.

In terms of the experimental equipment, we tried to have a "natural" looking robot, using clothes for the robot with neutral and not too bright colors. We tried to keep the experimental area 'tidy', i.e. tried not leave unrelated objects in the experimental area that could distract the attention from the main focus of the experiment.

Demos and explanations to participants were very important in our experiments; they were kept the same for all the participants. Giving slightly more explanation or a smile of the experimenter during the experiment might affect the evaluations of the participants. Also it is important to use the same experimenter for the all experiments since his/her behaviour or characteristics (gender, height, age, tone of voice etc.) may have an impact on participants in the experiments. Ideally experiments could be repeated with different experimenters, in order to reveal any effects this may cause, but practically this usually goes beyond the scope of an HRI study. In our HRI studies the experimenters had a 'passive role' during the experiments, i.e. they were present but did not proactively engage with the participants. The main role was to provide demos, explanations, and generally guide the participants through the experiment.

It is important to use small pilot groups to test the usability and understandability of the questionnaires before the real experiments. Especially when conducting large scale experiments, with strict time restrictions, this is a vital issue. In our experiment with children we could only use one child to test the questionnaire. Ideally several children should be used to get feedback in particular on the questionnaire design.

Surveys and questionnaires might not always provide the 'full picture'. For example, females tend to have higher agreeableness scores than males and participants with higher agreeableness can thus be expected to rate the robot's capabilities as better [Costa, et al., 2001]. Also if some of the participants who knew the experimenter, they might have be less 'subjective', or the primary school children might answer the questions in the same spirit as doing 'homework', and might not express their 'natural' feelings. So it is vital to collect additional behavioral data from different sources e.g. video recording, using different sensors, and comparing the behavioral data and questionnaire data for different aspects. Interestingly, our research [Kose-Bagci, et al., 2007; Kose-Bagci, et al., 2008] showed that results from different data sources might not 'agree'. For example, when judging the most preferred game, and least preferred game, although a participant might dislike the drumming of robot in one game and evaluate the game in the questionnaire according to this, he/she might have scored the least error in terms of objective measures of drumming performance in that game

It is important not to explain fully the overall experimental design and the game, not even details of the robot (in the embodiment experiments) to the participants before the start of the interaction experiments e.g. in instructions because the aim is to observe the adaptation of the human to the robot and different aspects of (non-verbal) communication emerging from the social interaction between the human participant and the humanoid. If the participant starts with too much information this may bias him/her and affect the results. Sometimes participants spend some time to explore the robot rather then playing the game, which increases performance error according to objective behavioural measures, but we did not intervene since we considered it a part of the interaction. We used the demo session to address this issue. Still participants asked questions at the end, for example: "Does the robot learn?", "Can it see me?", "Why does not it talk?" One of the participants claimed that the robot smiled at him 'badly' when he did something 'wrong' although the robot was not making any facial gestures at all.

The duration of the experiments is important. Our experiments with adults lasted 3 minutes and those with children 2 minutes. The experiment should be long enough to collect an adequate amount of data and short enough not to be boring since boredom also effects the evaluations of the participants. Even if the task is enjoyable and the humanoid is interesting, in experiments involving a repetition of movements or tasks, doing the same thing for several minutes is not always pleasant.

Although we worked with groups of participants at different times (these three experiments were completed over more than 1.5 years), we kept all the experimental conditions identical (e.g. the experimental setup and the robot gestures used), within the same experiment, which was a hard task when using a robot (which had been used extensively during the same period for several other research projects), and dealing with more than 100 human participants.

In all of our experiments, the robot was autonomous. Therefore it was important to avoid the belief in participants that the robot was remotely controlled by a Wizard-of-Oz (WoZ) technique, see e.g. [Green, et al., 2004]. During the trials we tried to avoid using the control laptop, because when the experimenter worked with the laptop, the participants might have thought that she was controlling the robot. Moreover, in the embodiment experiments, the children could think that not the robot but the experimenter was playing the drum in the disembodied condition (when the robot was hidden but its drumming sound was heard), so we always kept the experimenter in view of the participants, but not watching the participants, not using the laptop, and not interfering with the participant or the robot. Rather, the experimenter did something seemingly 'irrelevant' to the study, i.e. reading a book. Being watched may put stress on the participants.

The selection of the robot that was used in the HRI experiments is also an important issue. Some humanoid robot's are functional and robust from the experimenter's and designer's point of view, but might look 'scary' especially from children's point of view, compare research on the 'uncanny valley', e.g. [MacDorman & Ishiguro, 2005]. Some participants find the inner noise of the robot operating 'normal' as this makes it more 'robot-like', but it can be annoying for others.

Some people may have concerns towards robots which may prevent them from interacting with the robot 'naturally'. Such participants would tend to behave in a manner less 'relaxed' and 'open' towards the robot. Such an attitude might be hard to recognize in questionnaire data but can be detected in behavioral data. It is essential to gauge people's feelings regarding and attitudes towards robots in order to detect participants with strong negative feelings towards robots. Generally, participants' personality profiles, individual interests, hobbies etc, may also provide useful data that may explain how people react to and interact with robots in HRI experiments.

Related work

Other researchers have identified various important methodological issues in HRI research. However, a full survey of related work goes beyond the scope of this paper. Illustrating related work, Walters and colleagues (2005) also have provided a very useful discussion of the practical and methodological aspects of HRI studies which were based on several HRI experiments. They describe the legal and safety issues in detail. Those experiments took a human-centered perspective in HRI studies with a human-scaled mobile robot which was primarily controlled by the WoZ technique. The methodological issues related to these experiments, which are slightly different from those of our experiments, were described in detail. Table 1 shows a comparison of both works. Importantly, while Walters et al. (2005) used a primarily remote controlled robot, our experiments have taken a dual perspective: developing autonomous behaviours for a humanoid robot to play interaction games with people, while at the same time assessing the behaviour of people playing interaction games with the robot and their subjective evaluations of the games.

Table 1. Comparison of Drum-mate and Walters et al. Wizard-of-Oz (2005) studies

	Drum-mate studies	[Walters, et al., 2005] studies
Robot platform	Kaspar, a child sized humanoid	PeopleBot™, a human-sized wheeled robot base, extended
Mobility of the robot	Immobile, just head and arms move	Mobile
Control of the robot	autonomous	WoZ + (autonomously in a small scale)
Appearance	Human-like features	Mechanical looking
Experiments with participants	One participant at a time	One participant at a time and groups of participants
Experimenter	In view of participants to prove the robot is operating autonomously	Experimenters controlling the robot are hidden from the robot, the experimenter introducing the participants to the experiment etc. is in view of participants
Perspective of the trials	Both robot and human centered (development of interactive games for a humanoid robot, as well as the study of people's behaviour and subjective evaluation in interaction experiments)	Human centered (studying perceptions, attitudes and behaviour of people towards robots)

Conclusion

We presented the experimental design and the related issues result of three interaction experiments with an autonomous humanoid robot, involving in total 116 human participants playing human-humanoid drumming games. Despite the issues related to the experimental environment, participants, and the robot itself, we had significant results in terms of non-verbal and timing aspects of interaction, imitation, turn-taking and gender differences that are reported elsewhere [Kose-Bagci et al,. 2007; 2008; in preparation].

The methodological issues, and approaches taken to address these and other issues presented in this paper will inform future studies related to human-humanoid social interaction. Many of these issues will also play a role in other HRI experiments, including different application areas such as entertainment, service robots, and educational/therapy robots. Thus, we hope that this paper will be useful for other HRI researchers, in particular those with no or little experience in carrying out user studies.

References

Blow, M.P., K. Dautenhahn, A. Appleby, C. Nehaniv, D. Lee, "Perception of robot smiles and dimensions for human-robot interaction design*",* In *Proc. IEEE RO-MAN06,* pp.469-474, 2006.

Costa, PT; Terracciano, A; McCrae, RR. Gender differences in Personality traits across cultures: Robust and surprising findings. *Journal of Personality and Social Psychology*. 81:322–331, 2001.

FearNot! event, http://www.herts.ac.uk/news-and-events/latest-news/Antibullying-software-and-robots.cfm, Last accessed: 22 December, 2008.

Green, A., Huttenrauch, H., K. Severinson-Eklundh, "Applying the Wizard-of-Oz framework to cooperative service discovery and configuration", In *Proc. 13th IEEE international workshop on robot and human interactive communication (RoMan2004), Kurashiki, Japan*, 575–580, 2004.

Kose-Bagci, H., K. Dautenhahn, D. S. Syrdal, and C. L. Nehaniv, "Drum-mate: A Human-Humanoid Drumming Experience"*, IEEE-RAS Humanoids2007*, 2007.

Kose-Bagci, H., K. Dautenhahn, and C. L. Nehaniv, "Emergent Turn-Taking Dynamics in Drumming Games with a Humanoid Robot" , *Proc. IEEE RO-MAN 2008*, 1-3 Aug. 2008, Technische Universitat Munchen, Munich, Germany, pp. 346-353, 2008.

MacDorman, K. F. & Ishiguro, H., "The uncanny advantage of using androids in cognitive science research", *Interaction Studies, 7*(3), 297-337, 2006.

Walters, M. L., S Woods, K L Koay, & K Dautenhahn, Practical and methodological challenges in designing and conducting interaction studies with human subjects. *Proc. AISB'05 Symposium on Robot Companions Hard Problems and Open Challenges in Human-Robot Interaction*, 14-15 April 2005, University of Hertfordshire, UK, pp. 110-120, 2005.

Applied Bionics and Biomechanics
Vol. xx, No. xx, xx 2009, 1–29

Taylor & Francis
Taylor & Francis Group

SECOND PROOFS OF ARTICLE IN PRESS

KASPAR – a minimally expressive humanoid robot for human–robot interaction research

Kerstin Dautenhahn*, Chrystopher L. Nehaniv, Michael L. Walters, Ben Robins, Hatice Kose-Bagci, N. Assif Mirza and Mike Blow

School of Computer Science, University of Hertfordshire, Herts, United Kingdom

(*Received 17 October 2008; final version received 16 June 2009*)

This paper provides a comprehensive introduction to the design of the minimally expressive robot KASPAR, which is particularly suitable for human–robot interaction studies. A low-cost design with off-the-shelf components has been used in a novel design inspired from a multi-disciplinary viewpoint, including comics design and Japanese Noh theatre. The design rationale of the robot and its technical features are described in detail. Three research studies will be presented that have been using KASPAR extensively. Firstly, we present its application in robot-assisted play and therapy for children with autism. Secondly, we illustrate its use in human–robot interaction studies investigating the role of interaction kinesics and gestures. Lastly, we describe a study in the field of developmental robotics into computational architectures based on interaction histories for robot ontogeny. The three areas differ in the way as to how the robot is being operated and its role in social interaction scenarios. Each will be introduced briefly and examples of the results will be presented. Reflections on the specific design features of KASPAR that were important in these studies and lessons learnt from these studies concerning the design of humanoid robots for social interaction will also be discussed. An assessment of the robot in terms of utility of the design for human–robot interaction experiments concludes the paper.

Keywords: humanoid robots; minimally expressive robot; human–robot interaction; social interaction

1. Introduction

A key interest in our research group concerns human–robot interaction research; see Fong et al. (2003), Dautenhahn (2007), Goodrich and Schultz (2008) for introductory material of this research field. One of the most challenging open issues is how to design a robot that is suitable for human–robot interaction research, whereby suitability not only concerns the technical abilities and characteristics of the robot but, importantly, its perception by people who are interacting with it. Their acceptance of the robot and willingness to engage with the robot will not only fundamentally influence the outcome of human–robot interaction experiments but will also impact the acceptance of any robots designed for use in human society as companions or assistants (Dautenhahn et al. 2005; Dautenhahn 2007). Will people find a machine with a human appearance or the one that interacts in a human-like manner engaging or frightening? If a face is humanoid, what level of realism is suitable? Different studies have independently shown the impact of robot appearance on people's behaviour towards robots, expectation from and opinions about robots; see Walters (2008a) and Walters et al. (2008b) for in-depth discussions. Lessons learnt from the literature indicate that a humanoid appearance can support enjoyable and successful human–robot interaction; however, the degree of human-likeness required for a certain task/context etc. remains unclear.

In contrast to various approaches trying to build robots as visual copies of humans, so-called 'android' research (MacDorman and Ishiguro 2006), or research into designing versatile high-tech humanoid robots with dozens of degrees of freedom (DoFs) in movement and expression (cf. the iCub humanoid robot, Sandini et al. 2004), the approach we adopted is that of a humanoid, but minimally expressive, robot called KASPAR[1] that we built in 2005 and have modified and upgraded since then (see Figure 1). Our key aim was to build a robot that is suitable for different human–robot interaction studies. This paper describes the design and use of the robot.

In order to clarify concepts that are important to the research field of human–robot interaction, the following definitions of terms that are being employed frequently in this paper will be used[2]:

Socially interactive robots (Fong et al. 2003): Robots for which social interaction plays a key role, different from other robots in human–robot interaction that involve teleoperation scenarios.

*Corresponding author. Email: K.Dautenhahn@herts.ac.uk
All authors carried out the work while being part of the Adaptive Systems Research Group at University of Hertfordshire.

[1]KASPAR: Kinesics and Synchronization in Personal Assistant Robotics.

[2]Other related definitions relevant to the field of human–robot interaction and social robotics are discussed in Dautenhahn (2007).

ISSN: 1176-2322 print / 1754-2103 online

DOI: 10.1080/11762320903123567
http://www.informaworld.com

Figure 1. The minimally expressive humanoid robot KASPAR designed for social interaction.

Humanoid robots, humanoids (based on Gong and Nass 2007): "A robot which is not realistically human-like in appearance and is readily perceived as a robot by human interactants. However, it will possess some human-like features, which are usually stylised, simplified or cartoon-like versions of the human equivalents, including some or all of the following: a head, facial features, eyes, ears, eyebrows, arms, hands and legs. It may have wheels for locomotion or use legs for walking" (Walters et al. 2008b, p. 164).

Of specific interest to the present paper are humanoid robots with faces. Generally these can range from abstract/cartoon-like to near-to-realistic human-like faces. Section 2.2.2 discusses in more detail the design space of robot faces and Section 3 motivates our decision for a *minimally expressive* face.

This paper has been structured as follows: Section 2 provides an introduction to important issues in the design of robots and robot faces, in particular with respect to the design space of robots and how people perceive and respond to faces. Related work and design issues discussed in the literature are critically reflected upon. Section 3 describes the issues and rationale behind the design of minimally expressive humanoids in general and KASPAR in particular, and provides construction details regarding the current versions of the robot used in research. Section 4 illustrates its use in a variety of projects covering the spectrum from basic research to more application-oriented research in assistive technology. Human–robot interaction studies with KASPAR are summarised and discussed in the light of KASPAR's design features. The conclusion (Section 5) reflects upon our achievements and provides a conceptual assessment of KASPAR's strengths and weaknesses.

2. Robot design for interaction

This section reflects in more detail on issues regarding the appearance of a robot in the context of human–robot interaction and how people perceive faces (robotic or human). Related work on designing socially interactive research platforms will be discussed. Note, we do not discuss in detail the design of commercially available robots since usually little or nothing is made public about the details or rationale of the design. An example of such robots is the Wakamaru (Mitsubishi Heavy Industries), which has been designed to 'live with humans'. Unfortunately only brief, online information has been provided about the design rationale, hinting at the importance of expressiveness in the eyes, mouth and eyebrows (Wakamaru 2009).

Thus, for a more detailed comparison of the design rationale of KASPAR with other robots, we focus our discussion of related work on other *research* platforms.

2.1. The design space of humanoid robots

The effect of the aesthetic design of a robot is an area that has often been neglected, and only in visual science fiction media or recently with the advent of commercial household robots has it been paid much attention. A notable exception is the 'uncanny valley' proposed by Masahiro Mori (Mori

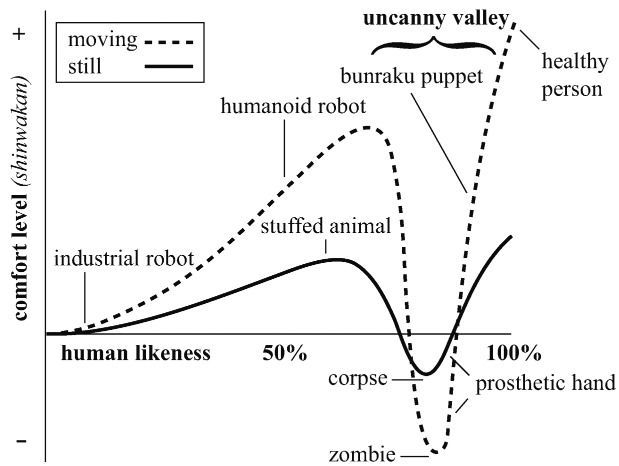

Figure 2. The uncanny valley (MacDorman et al. 2009).

1970). Mori proposed that the acceptance of a humanoid robot increases as realism increases, up to a point where, as the robot approaches perfect realism, the effect becomes instead very disturbing and acceptance decreases sharply because the robot starts to look not quite human or at worst like a moving corpse (see Figure 2 to illustrate the 'uncanny valley'). In theory the realism of both appearance and movement can give rise to this effect, with movement evoking the stronger response. It is possible that there may also be 'behavioural uncanniness' affecting perception of a robot during social interaction and governed by (among other things) the appropriateness and timing of its responses to social cues. However, little empirical data exists to support Mori's theory and opinions vary as to the strength of the effect and its longevity; see MacDorman (2005a) and MacDorman et al. (2005b) for recent works on the uncanny valley.

Previous work has identified a number of issues that are important in the design of robots meant to socially interact with people (Dautenhahn 2004). A full review of the technical and theoretical aspects of different robot designs in the field of humanoid robotics would go beyond the scope of this paper; therefore we discuss in the following paragraphs in more detail the key design features of the robot Kismet. Both Kismet and KASPAR have been specifically designed for human–robot interaction and, importantly, detailed information about the design rationale of Kismet is available in the research literature.

When Breazeal (2002) designed Kismet, which '. . . is designed to have an infant-like appearance of a fanciful robotic creature' (p. 51), with a youthful and appealing appearance, her intention was not to rival but rather to connect to the social competence of people. Furthermore, she incorporated key features in the robot that are known to elicit nurturing responses, as well as other non-humanoid features (e.g. articulated ears), in conjunction with exaggerated, cartoon-like, believable expressions. The overall cartoon-like appearance of the robot took advantage of people's liking and familiarity with cartoon characters. The overall design has been very successful: 'As a result, people tend to intuitively treat Kismet as a very young creature and modify their behavior in characteristic baby-directed ways' (Breazeal 2002, p. 51). However, it should be noted that the robot has never been used in any task-oriented scenarios that involve the manipulation of objects due to the fact that it does not have any manipulation abilities. The overall design is based on the assumption that people *are eager* to interact with the robot in the role of a caretaker. We contend that while this may be an appropriate approach for entertainment purposes, it is unclear how this design approach of a 'robotic pet/baby' would apply to work that is oriented towards robots as assistants or companions (see a detailed discussion of these two different approaches in Dautenhahn 2007). Note, Kismet was an expensive laboratory prototype, and to run its sophisticated perception and control software required more than 10 networked PCs.

In Breazeal and Foerst (1999) several of Kismet's design guidelines are presented for achieving human–infant like interactions with a humanoid robot; however, the underlying basic assumption here is 'the human as a caretaker', so some, but not all, of these guidelines are relevant for this paper. We now discuss these guidelines in relation to the specific approach that we took with the design of our humanoid robot KASPAR.

Issue I: 'The robot should have a cute face to trigger the 'baby-scheme' and motivate people to interact with it, to treat it like an infant, and to modify their own behavior to play the role of the caregiver (e.g. using motherese, exaggerated expressions and gestures)'.

Cuteness of the robot is not a key issue in the design rationale of our robot KASPAR because we did not envisage human–infant caretaker interactions. On the contrary, our goal was to have a robot that people may relate to in different ways, depending on the particular context of use and application domain.

Issue II: 'The robot's face needs several degrees of freedom to have a variety of different expressions, which must be understood by most people. Its sensing modalities should allow a person to interact with it using natural communication channels'.

Our approach partly agrees with the view on this issue; however, we focused on what we call a *minimally expressive face* with few expressions and few sensors in order to emphasise the most *salient* human-like cues of the robot. Rather than trying to make a robot very human-like, our goal was to concentrate on a few salient behaviours, gestures and facial expressions in order to run experiments that systematically study the influence of each of these cues on the interaction with people. Note, while Kismet also includes some cues that are zoomorphic but not anthropomorphic (e.g. articulated ears), the design of KASPAR's face focused on human-like features alone in order not to violate the aesthetic consistency.

Issue III: 'The robot should be pre-programmed with the basic behavioral and proto-social responses of infants. This includes giving the robot the ability to dynamically engage a human in social [interaction]. Specifically, the robot must be able to engage a human in proto-dialogue exchanges'.

Our approach uses an emphasis on non-verbal interaction without any explicit verbal 'dialogue'. We are interested in the *emergence* of gesture communication from human–robot interaction dynamics. Also, rather than solely building a research prototype for the laboratory, our aim was to have a robot that can be used in different application areas, including its use in schools, and under different methods of control (remote control of the robot as well as autonomous behaviour).

Issue IV: 'The robot must convey intentionality to bootstrap meaningful social exchanges with the human. If the human can perceive the robot as a being "like-me", the human can apply her social understanding of others to predict and explain the robot's behavior. This imposes social constraints upon the caregiver, which encourages her to respond to the robot in a consistent manner. The consistency of these exchanges allows the human to learn how to better predict and influence the robot's behavior, and it allows the robot to learn how to better predict and influence the human's behavior'.

The above is again very specific to the infant–caretaker relationship that Kismet's design is based upon. Rather than a 'like-me' perception of the robot we targeted a design that allows a variety of interpretations of character and personality on the robot (which might be termed 'it could be me' – see Dautenhahn 1997). Below we discuss this issue in more detail in the context of the design space of faces.

Issue V: 'The robot needs regulatory responses so that it can avoid interactions that are either too intense or not intense enough. The robot should be able to work with the human to mutually regulate the intensity of interaction so that it is appropriate for the robot at all times'.

Issue VI: 'The robot must be programmed with a set of learning mechanisms that allow it to acquire more sophisticated social skills as it interacts with its caregiver'.

Issues V and VI as discussed by Breazeal and Foerst relate specifically to the programming of the robot. For KASPAR we did not aim at a 'pre-programmed' robot but intended to build an open platform that would allow the development of a variety of different controllers and algorithms.

Other related work on humanoid robots includes the Lego robot Feelix (Cañamero 2002) that reacts to tactile stimulation by changing its facial expression. Feelix follows a similar design approach as Kismet, i.e. using exaggerated features, but a low-cost approach with commercially available Lego components. The humanoid robot Robota (Billard et al. 2006) has been designed as a toy for children

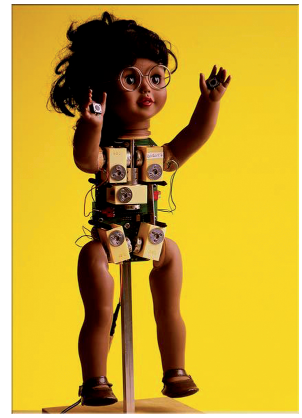

Figure 3. Robota (Billard et al. 2006).

and used in various projects involving imitation, interaction and assistive technology (Robins et al. 2004a, 2004b, 2005b). The key movements of this robot in these studies include turning of the head (left and right movements) and lifting of arms and legs (up and down movements of the whole limbs). Facial expressiveness or the generation of more complex gestures was not possible. The design considerations of Robota (see Figure 3) addressed in (Billard et al. 2006) include the following:

1. *Ease of Set-up*: This concerns the ease of setting up sessions, e.g., in schools, and favours a light-weight, small-sized and low-cost robot with on-board processing and battery power.

Note, the above design consideration applies generally to all robots that are meant to be used in different locations where they have to be brought 'in and out' quickly, different from a robot that relies on a sophisticated laboratory set-up (such as above-mentioned Kismet). Since the robot whose design we were undertaking was also meant to be suitable for school applications, it was important for us, too, to keep the costs down. We decided that the price of the robot should be comparable to that of a laptop.

2. *Appearance and behaviour*: This criterion concerns the human-likeness in the appearance of the robot. Robota had a static face (from a toy doll), so it included some human-like features. A doll-like appearance was also considered to be 'child-friendly'. Billard et al. (2006) argued that taking a doll as a basis would help to integrate the robot in natural play environments.

The above design considerations are consistent with our approach to the design of KASPAR, where we used a mannequin as the basis of the 'body' of the robot; however, we replaced the head (including the neck) and designed a minimally expressive robot. Thus, while the design of KASPAR started before Billard et al.'s publication of design guidelines (2006), several key aspects are common.

Other research groups have studied the design of robots for 'child's play', including Michaud et al. (2003) who discuss design guidelines for children with autism but with an emphasis on mobile robots and playful interactions as related to the robot's behaviour, focusing primarily on non-humanoid robots. This work indicates that the design space of robots is vast, and, depending on the actual user groups and requirements as well as individual needs and preferences, different designs may be favourable. Different from this work, in the context of this paper we focus on minimally expressive humanoid robots, suitable for human–robot interaction experiments in assistive technology as well as developmental robotics research. Please note, in Section 4.1 we discuss in more detail design issues of robots for the particular application area of autism therapy.

Since the key component of KASPAR is its minimally expressive face and head, the next sections provide more background information on the perception of faces.

2.2. Perceptions of faces

In this section we discuss some important issues to how people perceive human or robot faces.

2.2.1. Managing perceptions

DiSalvo et al. (2002) performed a study into how facial features and dimensions affect the perception of robot heads as human-like. Factors that increased the perceived humanness of a robot head were a 'portrait' aspect ratio (i.e. the head is taller than its width), the presence of multiple facial features and, specifically, the presence of nose, mouth and eyelids. Heads with a 'landscape' aspect ratio and minimal features were seen as robotic. They suggest that robot head design should balance three considerations: 'human-ness' (for intuitive social interaction), 'robot-ness' (to manage expectations of the robot's cognitive abilities) and 'productness' (so that the human sees the robot as an appliance). The idea of designing a robot to be perceived as a consumer item is noteworthy for the fact that people's *a priori* knowledge of electronic appliances can be utilised in avoiding the uncanny valley; the implication is that the robot is non-threatening and under the user's control. To fulfil their design criteria, they present six suggestions: a robot should have a wide head, features that dominate the face, detailed eyes, four or more features, skin or some kind of covering and an organic, curved form.

2.2.2. The design space of faces

Faces help humans to communicate, regulate interaction, display (or betray) our emotions, elicit protective instincts, attract others and give clues about our health or age. Several studies have been carried out into the attractiveness of human faces, suggesting that symmetry, youthfulness and skin condition (Jones et al. 2004) are all important factors. Famously, Langlois and Roggman (1990) proposed that an average face – that is, a composite face made up of the arithmetic mean of several individuals' features – is fundamentally and maximally attractive (although there are claims to the contrary, see Perrett et al. 1994), and that attractiveness has a social effect on the way we judge and treat others (Langlois et al. 2000).

Human infants seem to have a preference for faces, and it appears that even newborns possess an 'innate' ability to spot basic facial features, such as a pair of round blobs situated over a horizontal line which is characteristic of two eyes located above a mouth. It has been debated whether this is due to special face recognition capability or due to sensory-based preferences for general perceptual features such as broad visual cues and properties of figures such as symmetry, rounded contours etc., which then, in turn, form the basis for learning to recognise faces (Johnson and Morton 1991). The nature and development of face recognition in humans is still controversial. Interestingly, while the baby develops, its preference for certain perceptual features changes until a system develops that allows it to rapidly recognise familiar human faces. Evidence suggests that exposure to faces in the first few years of life provides the necessary input to the developing face recognition system (see Pascalis et al. 2005). The specific nature of the face stimuli during the first year of life appears to impact the development of the face-processing system. While young infants (up to about six months of age) can discriminate among a variety of faces belonging to different species or races, children at around nine months (and likewise adults) demonstrate a face-representation system that has become more restricted to familiar faces. The social environment, i.e. the 'kinds of faces' an infant is exposed to, influences the child's preferences for certain faces and abilities to discriminate among them. Not only time of exposure but also other factors, including emotional saliency, are likely to influence the tuning of the face recognition systems towards more precision (Pascalis et al. 2005).

In terms of perception of emotions based on faces, it is interesting to note that people can perceive a variety emotions based on rigid and static displays, as exemplified in the perception of Noh masks that are used in traditional Japanese Noh theatre. Slight changes in the position of the head of an actor wearing such a mask lead to different types of emotional expressions as perceived by the audience. This effect is due to the specific design of the masks where changes in angle and lighting seemingly 'animate' the face. Lyons et al. (2000) scientifically studied this effect (see Figure 4) and also pointed out cultural differences when studying Japanese as well as British participants. We are not aware if this Noh mask effect has

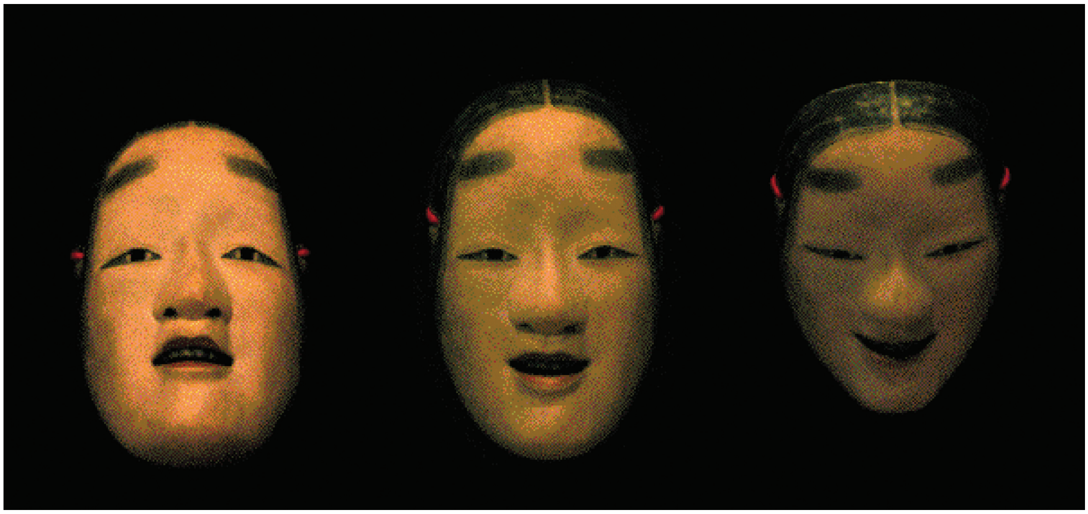

Figure 4. The Noh mask effect. Photo used with permission (Lyons et al. 2000).

been exploited deliberately in the design of humanoid robot expressions.

In his book *Understanding Comics* (McCloud 1993) on narrative art, Scott McCloud introduces a triangular design space for cartoon faces (Figure 5). The left apex is realistic, i.e. a perfect representation of reality, for example a photograph, or realistic art such as that by Ingres. Travelling to the right faces becomes more iconic, that is, the details of the face are stripped away to emphasise the expressive features; emoticons such as ':)' are a perfect example in the 21st century zeitgeist. The simplification has two effects. Firstly, it allows us to amplify the meaning of the face, and to concentrate on the message rather than the medium. Secondly, the more iconic a face appears the more people it can represent. Dautenhahn (2002) points out that iconography can aid the believability of a cartoon character. We are more likely to identify with Charlie Brown than we are with Marilyn Monroe, as a realistic or known face can only represent a limited set of people, whereas the iconic representation has a much broader range – to the extent of allowing us to project some aspects of ourselves onto the character. Towards the top apex representations become *abstract*, where the focus of attention moves from the meaning of the representation to the representation itself. Examples in art would be (to a degree) Picasso's cubist portraits or the art of Mondrian.

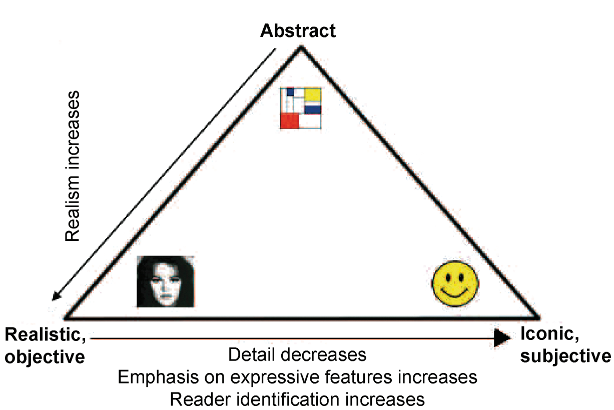

Figure 5. The design space of comics (Blow et al. 2006), modified from McCloud (1993). Note, similar principles are also relevant to animation and cartoons.

We can use this design space, and the accumulated knowledge of comic's artists, to inform the appearance of our robots. Figure 6 shows some robot faces and their (subjective) places on the design triangle. Most are 'real-life' robots although several fictional robots have been included, as functionality has no bearing on our classification in this context. At the three extremes are NEC's Papero (iconic), a small companion robot which is relatively simple and cheap to make and allows easy user-identification; Hanson's K-bot (realistic), complex and theoretically deep in the uncanny valley but allowing a large amount of expressive feedback and Dalek (abstract), potentially difficult to identify with but not as susceptible to the uncanny valley due to its non-human appearance.

Of course, the design space only addresses the static appearance of the robot. The nature of most robot faces is that they encompass a set of temporal behaviours that greatly affect our perception of them. For example, as these issues are so important in human–human interaction (Hall 1983), it seems well worthwhile investigating the rhythm and timing of verbal and, especially, non-verbal behavioural interaction and dynamics of robots interacting with humans, an area referred to as *interaction kinesics* (Robins et al. 2005a). An extension of McCloud's design space to investigate

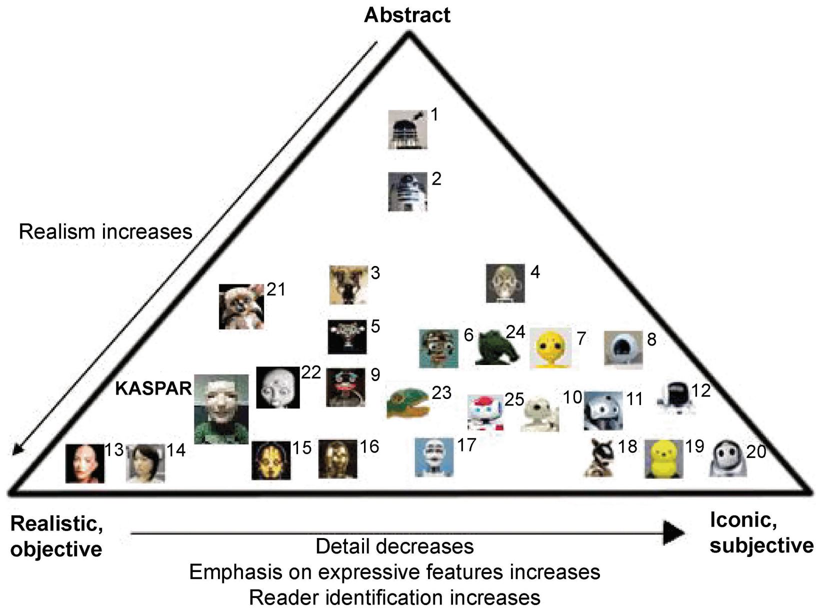

Figure 6. Robot faces mapped into McCloud's design space, updated version of Blow et al. (2006). (1) Dalek (© the British Broadcasting Corporation/Terry Nation); (2) R2D2, fictional robot from 'Star Wars' (© Lucas Film Ltd.); (3) DB (© ATR Institute Kyoto); (4) MIT humanoid face project (© MIT); (5) Kismet (© MIT/Cynthia Breazeal); (6) Infanoid (© Hideki Kozima); (7) Wakamaru communication robot (© Mitsubishi Heavy Industries); (8) HOAP-2 (© Fujitsu Automation); (9) Minerva tour-guide robot (© Carnegie Mellon University); (10) Toshiba partner robot (© Toshiba); (11) QRIO (© Sony); (12) ASIMO (© Honda); (13) K-Bot, extremely realistic 24 DoF head built by David Hanson (© Human Emulation Robotics); (14) Repliee-Q1 (© Osaka University/Kokoro Inc.); (15) False Maria, fictional robot from Fritz Lang's 1927 film 'Metropolis'; (16) C3PO, fictional robot from 'Star Wars' (© Lucas Film Ltd.); (17) WE-4R robot (© WASEDA University); (18) AIBO robotic dog (© Sony); (19) Keepon, minimal DoF HRI robot (© Hideki Kozima); (20) Papero household robot (© NEC); (21) Leonardo HRI research robot (© MIT Personal Robots Group); (22) Nexi HRI research robot (© MIT Personal Robots Group); (23) Pleo commercial companion robot (© Ugobe Inc.); (24) Probo medical companion robot for children (© Vrije Universiteit Brussel); (25) Nao personal robot (© Aldebaran Robotics).

behavioural aspects would be a worthwhile study, specifically how a robot's behaviour affects its perception as iconic, realistic or abstract, and the effect of social behaviour on the uncanny valley and user identification with the robot.

As one moves in the design space of faces from realism towards iconicity, a human is more likely to identify themselves with the face due to the decrease in specific features, and the distinction between *other* and *self* becomes less and less pronounced. Could this idea be useful in robot design? If a robot is to be designed to extend humans' abilities or carry out tasks on their behalf, iconic features may possibly allow the user to project their own identity onto the robot more easily. In contrast, realistic face designs will be seen objectively as *someone else*, and abstract designs often as *something else*. In this case the interaction partner's identification with the robot will be discouraged by the non-iconic nature of the design. Some robot roles (such as security guards) might benefit from reinforcing this perception. While the idea of the robot as an extension of self remains speculative at this point, future work in this area needs to shed more light on these issues.

3. Design of KASPAR

This section details the technical design of KASPAR. We start with general considerations for the design-space of minimal expressive humanoids, particular initial design requirements for KASPAR and then present the technical design and construction details.

3.1. Robot design and construction details

3.1.1. General considerations for the design-space of minimal expressive humanoids

First we discuss some key considerations on the expressive face/head and general appearance and expression in minimal expressive humanoids for human–robot social interaction. In the next section the requirements for KASPAR are introduced.

3.1.2. Balanced design

- If face, body and hands are of very different complexities, this might create an unpleasant impression for humans interacting with the robot. Aesthetic coherence also requires balance in the physical design and

in turn also the behavioural and interactional design of the robot and its control.
- DoFs and design should be appropriate for the actual capabilities that the robot will possess and use (otherwise inappropriate expectations are created in the human). (cf. Dautenhahn and Nehaniv 2000).

3.1.3. Expressive features for creating the impression of autonomy

- Attention: Visible changes in direction of head, neck and eye gaze direction (i.e. with independent DoFs within eyes) are the most important expressive features in creating the impression of autonomy. In a humanoid, this entails actuation of the neck in some combination of pan, tilt and roll.
- Emotional state: Expressive components in face (eyes, eyebrows, mouth, possibly others) are at the next level of importance (see point 3 below).
- Contingency: The human interaction partner should see *contingency* of the robot's attentional and expressive states as it responds to interaction – this entails behavioural design on appropriate hardware (see below for minimal 6+ DoF systems).

Conveying attention (indication of arousal and direction of attention) and the impression of autonomy has been illustrated in the elegant design of the very minimal, non-humanoid robot Keepon by Hideki Kozima (Kozima et al. 2005).

3.1.4. Minimal facial expressive features

One can make use of the Noh mask-like effects discussed above. This may be compared to Y. Miyake's concept of *co-creation* in man–machine interaction, namely that a human's subjective experience of a technological artifact, such as a robot or *karakuri* (traditional Japanese clockwork automaton), lies in the situated real-time interaction between observer, artifact and the environmental situation (Miyake 2003; see also Dautenhahn 1999). Therefore, we propose that a largely still, mask-like face (or even other body parts) that is dynamically oriented and tilted at different angles can be designed and used to induce various perceptions of the robot's state in the interaction with a human participant.

Unlike extreme minimal robots (such as Keepon) or robots with complex facial actuation expressiveness in the head (e.g. Kismet) in conjunction with the Noh-like elements of design, a few DoFs within the head may provide additional expressiveness (e.g. smiling, blinking, frowning, mouth movement etc.). Human-like robots with such minimal degrees of face actuation include Feelix by Lola Cañamero at University of Hertfordshire (Cañamero 2002) and Mertz by Lijin Aryananda at MIT-CSAIL (Aryananda 2004).

Possibilities for this additional facial actuation (approximately 6+ DoFs) are included:

- Eyebrows: 270° rotary 1 DoF/eyebrow (× 2), RC servo; if an additional DoF is to be used, then it could be used for raising/lowering the eyebrow in the vertical direction. (Eventually, directly actuated eyebrows were dropped from the first design of KASPAR in order to maintain aesthetic coherence. The adopted design leads to indirect expressiveness via the eyebrows of the face mask under deformations due to mouth and smile actuation.)
- Eyes: Pan and tilt, possibly supporting mutual gaze and joint attention.
- Eyelids: Blinking (full or partial, at various rates).
- Lips/mouth: Actuators for lips to change shape of mouth, e.g. from horizon lips to open mouth, possibly more DoFs a right and left edge to lift/lower mouth (smile/frown); also opening/closing of mouth.

In a minimally expressive robot, some subset of the above features could be selected (e.g. direct actuation of emotional expression could be omitted completely, while retaining the capacity to show direct attention, or, if included, any combination of eyebrows, eyelids or mouth actuation etc., could be omitted).[3]

3.1.5. Specific requirements for a minimally expressive humanoid suitable for different human–robot interaction studies: KASPAR

The overall minimally expressive facial expressions of KASPAR have been designed in order not to 'overwhelm' the observer/interaction partner with social cues but to allow him or her to individually interpret the expressions as 'happy', 'neutral', 'surprised' etc. Thus, only as few motors were used that were absolutely necessary to produce certain *salient* features.

- Similar to Kismet, as discussed above, KASPAR was meant to have a youthful and aesthetically pleasing design. Different from Kismet, we did not want to explicitly elicit nurturing responses in people, but instead support the function of KASPAR as a playmate or companion. So we refrained from exaggerated facial features and decided on a *minimally* expressive face.
- It was considered important that the robot has the size of a small child, in order not to appear threatening.
    - KASPAR sits on a table in a relaxed playful way with the legs bent towards each other (the way children often sit when playing).

[3] We thank H. Kozima for discussions on the design of Keepon and A. Edsinger-Gonzales for technical discussions on the implementation of Mertz.

  - The head is slightly larger in proportion to the rest of the body, inspired by comic's design as discussed above (in order not to appear threatening).

- Unlike Kismet which requires a suite of computers to run its software, we decided to have KASPAR's software running either on-board the robot or from a laptop. The reason for this was that we envisaged KASPAR to be used in various human–robot interaction studies, including studies outside the laboratory, so the robot had to be easily transportable, easy to set up etc.
- A low-cost approach was also considered practical in case future research or commercial versions were planned (e.g. to use KASPAR as a toy, or educational/therapeutic tool in schools or at home).
- In order to have a 'natural' shape, a child-sized mannequin was used as a basis. The legs, torso and the hands were kept. The hands were not replaced by articulated fingers in order to keep the design simple, and in order to invite children to touch the hands (which is more like touching a doll).
- Arms were considered necessary for the study of gesture communication, and they also allow the manipulation of objects which is important for task-based scenarios, e.g. those inspired by children's play. It was decided to build low-cost arms with off-the-shelf components that are not very robust and do not allow precise trajectory planning etc, but can nevertheless be 'powerful' in interaction for producing gestures such as waving, peek-a-boo etc.
- The neck was designed to allow a large variety of movements, not only nodding and shaking the head, but also socially powerful movements such as slightly tilting the head (important for expressing more subtle emotions/personality traits such as shyness, cheekiness etc.).
- KASPAR has eyelids that can open and close. Blinking can provide important cues in human–human interaction, so we decided that this was a salient feature to be added.

3.1.6. Technical design considerations

A main criteria for KASPAR emphasised the desirability of low cost. The budget for KASPAR allowed up to 1500 Euros for material costs. Therefore, the following decisions were made at the initial design specification stage:

A shop window dummy modelled after an approximately two-year-old girl was available at reasonable cost. It already possessed the overall shape and texture required for the body of the robot and could be readily adapted to provide the mainframe and enclosure for the robot system's components. Therefore, it was decided that KASPAR would be stationary and would not have moving or articulated legs.

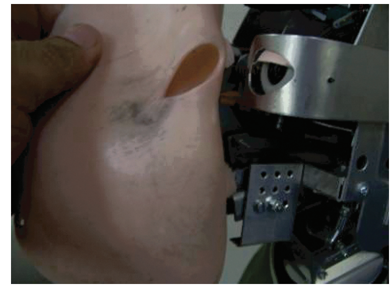

Figure 7. Detailed view of face mask attachment points.

In line with our discussion of identification and projection (as for Noh masks), it was also decided that the silicon rubber face mask from a child resuscitation practice dummy would be used for the face of the robot. These masks were flesh coloured and readily available as spare parts (to facilitate hygienic operation of the dummy). The masks were also sufficiently flexible to be deformed by suitable actuators to provide the simple expression capabilities that would be required, and also provided simplified human features which did not exhibit an unnerving appearance while static (cf. the 'uncanny valley' mentioned above, Mori 1970). See Figure 7 for the attachment of the mask to the robot's head.

It was decided that all joint actuation would be achieved by using radio control (RC) model servos. These were originally made for actuating RC models, but as they have been commercially available to the mass hobby market at low cost, they are also commonly used as joint actuators for small-scale robots. Interface boards are also available which allow them to be interfaced and controlled by a computer.

The main moving parts of the robot were head, neck and arms; hence, the original head, neck and arms were removed from the shop dummy to allow replacement with the respective new robot systems. The batteries and power and control components were fitted internally. KASPAR's main systems are described in more detail in the following sub-sections.

Further details of the design and construction of head and arms, as well as details of the robot's control and power supply are provided in Appendix A.

3.1.7. KASPAR II

About a year after completing KASPAR we built a second version called 'KASPAR II', and both robots are currently used extensively in different research projects. KASPAR II had been used in experiments on learning and interaction histories as reported in Section 4.3 (all other studies mentioned in this paper used the original KASPAR robot). KASPAR II's design is very similar to the original (KASPAR I), with a few modifications primarily in terms of upgrades. Details of KASPAR II are given in Appendix B, which also provides information on upgrades, changes and planned future improvements of KASPAR.

3.1.8. Remote control of the robot

In applications involving children with autism (see Section 4.1), a remote control was used to operate KASPAR. It is made of a standard wireless keypad (size 8 cm × 12 cm) with 20 keys. Different keys were programmed to activate different behaviours in KASPAR, i.e left/right arm drumming, waving, different postures etc. These are dynamic expressive behaviours released via single key press. The programmed keys had stickers on them with simple drawings representing the behaviour, e.g. a drum for drumming (two keys – right and left), a smiley for a 'happy' posture, a hand for hand-waving etc. The remote control allowed the introduction of collaborative games and role switch, with a view to use the robot as a social mediator, as will be explained in more detail in Section 4.1.

3.2. *Software*

The software development of KASPAR is not the focus of this paper and will thus only be mentioned briefly. The robot can be used in two modes: remotely controlled as well as autonomous operation. Unskilled operators can easily run and develop programmes for the robot using the novel user-friendly KWOZ (KASPAR Wizard of OZ) graphic user interface (GUI) software which runs on any Windows or Linux PC. This interface has been used in human–robot interaction scenarios when an experimenter (usually hidden from the participants) remotely controlled the robot from a laptop (see Section 4.1). This type of control is different from the remote control device that was specifically introduced to openly introduce collaborative games (see Section 3.1.8).

In a variety of projects KASPAR operates autonomously, see examples in Sections 4.2 and 4.3. An applications programming interface (API) provides access for programmers to develop custom programmes and access to open source robot software produced under the Yet Another Robot Programme (YARP) initiative (Yarp 2008).

3.3. *Aesthetics of the face*

As mentioned above, a child resuscitation mask was used.[4] The mask is produced by the Norwegian company Laerdal, which specialises in medical simulators and first produced 'Resusci-Anne', as a life-like training aid for mouth-to-mouth ventilation. Anne's face mask had been inspired by the 'peaceful-looking and yet mysterious death mask' (Laerdal Products Catalogue 2008–2009) of a girl who is said to have drowned herself in the Seine. The death mask is said to have first appeared in modellers' shops in Paris around the 1880s. In a 1926 catalogue of death masks it is called 'L'Inconnue de la Seine' (the unknown woman of the Seine). Replicas of the mask became fashionable as a decorative item in France and Germany. The mask and as yet unconfirmed stories surrounding its origin then sparked the imagination of many poets and other artists, such as Rilke, for the next few decades and led to numerous literary art works (The Guardian Weekend, 2007). The mysterious and beautiful, 'timeless' quality of the mask may contribute to its appeal to participants in human–robot interaction studies. In our view, the mask itself has a 'neutral expression' in terms of gender as well as age. It has a skin colour, without facial hair or any additional colouring, and we left it unchanged in order to allow viewers/interaction partners to impose different interpretations of personality/gender etc. on the robot.

Interestingly, the specific design and material that the rubber mask is made of, in conjunction with the attachment of the mask to the actuators, creates KASPAR's unique smile, which is minimal but naturalistic and similar to the so-called 'genuine smile' or 'true smile' shown by people. Ekman et al. (1990) describe the Duchenne smile (the genuine smile) that is characterised by movements of the muscles around the mouth and also the eyes. Humans show a true smile typically involuntarily. This smile is perceived as pleasant and has positive emotions associated to it, in contrast to other smiles in which the muscle orbiting the eye is not active. A variety of other smiles can be observed and they occur, e.g. when people voluntarily try to conceal negative experience (masking smiles), feign enjoyment (false smiles) or signal that they are willing to endure a negative situation (miserable smiles).

KASPAR's smile causes a very slight change in the mask around the eyes. This change is based on passive forces pulling on the mask when the mouth moves. Thus, this 'true' smile is possible due to particular way in which the smile was designed, how the mask is attached and the material properties of the mask.

As a consequence, KASPAR's smile is very appealing (Figure 8), and similar to a genuine smile shown by people. This is a novel feature that is different from many other robot (head) designs where smiles often appear 'false' since they either only operate the mouth or different parts of the face but not in the naturally smooth and dynamic fashion it occurs in KASPAR's face mask.

Note that the *dynamic* transition of the facial expressions (i.e. from neutral to a smile, cf. Figure 8) plays an important part in how people perceive KASPAR's facial expressions. Experimental results of an online survey with 51 participants (Blow et al. 2006) have shown that natural transitions (taking about two seconds from neutral expression to smile) are seen as more appealing than sudden (artificially created) transitions. Also, the larger the smile, the greater the participants' judgement of 'happiness'. However, while smiles with a natural transition are seen as more appealing than static pictures of the smiles, those with a sudden

[4]Thanks to Guillaume Alinier of the Hertfordshire Intensive Care & Emergency Simulation Centre at University of Hertfordshire for his generous donation of the face mask.

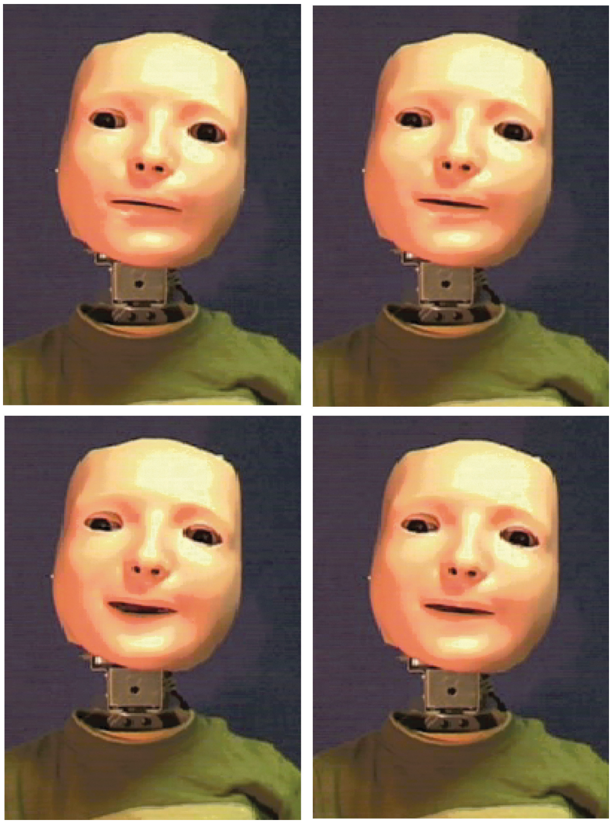

Figure 8. KASPAR's minimally expressive face illustrating four expressions designed for human–robot interaction. Clockwise from top left: neutral, small, medium and large smiles.

transition are not (Blow et al. 2006). This emphasises the need for consistency of appearance (in this case a humanoid face with a natural smile) and behaviour (the transition time of facial expressions). Further results of this study show that all four of KASPAR's expressions (Figure 8) shown to the participants were found appealing or very appealing. Note, our primary research interest is in human–robot interaction, not in facial design or emotion modelling, but these results give encouragement to participants' ratings of KASPAR's facial expressions. Other researchers might use KASPAR for a further investigation of these issues concerning the perception of robot facial expressions.

3.4. Contextual features

Contextual features are an important ingredient of interaction design (Preece et al. 2002). In order to help people relate to the robot socially we used various contextual features in terms of the robot's clothing. We dressed the robot in children's clothing (shirt, trousers and socks). We utilised children's used clothing which appear more natural than newly purchased clothing. We did not try to hide the fact that KASPAR is a robot, on the contrary we left the neck and wrists uncovered so that cables and pieces of metal can be seen.

For the applications of the robot in autism therapy (see Section 4.1) where we mainly work with boys, we wanted to give the robot a boyish appearance and added a baseball cap and a wig in order to emphasise the child-sized and playful nature of the robot. We tried different hair colours, but the dark-coloured wig gave the most consistent appearance. The cap can also serve as a prop and invites children to remove and replace it etc. Moreover, in several research projects where we study human–humanoid interaction games, we place a toy tambourine in the robot's lap, which the robot is able to drum on. This feature adds to the robot's perceived playfulness and allows the study of task-based interaction (e.g. drumming).

3.5. Gestures

As discussed above, our initial requirements were to have arms that allow simple gestures. During the course of using KASPAR in different research projects a number of dynamic gestural expressions were defined (Figure 9).

Note, while within our human–robot interaction research group we did not systematically study how different user groups perceive KASPAR's appearance and behaviours, we have been using the robot in multiple experiments, demonstration and public engagement events involving children and adults of different age ranges, gender, background etc. In total, more than 600 children have been exposed to the robot (either watching its live demonstrations or participating in an interaction experiment) along with about 300 adults. These encounters were part of interaction experiments carried out in schools or in the laboratory, or were part of public engagement events taking place either in schools, museums or conference venues, or on university premises. While feedback from the public events was very informal in nature, we nevertheless have gained anecdotal evidence that can be described as follows:

- Children of various ages (typically developing children as well as children with special needs, including children with autism; cf. Section 4.1) generally show a very positive reaction towards KASPAR, attempting spontaneously to play and interact with the robot, often touching it etc. The minimal facial expressions and gestures appear particularly appealing, the child-like appearance and size of the robot seems to elicit play behaviour similar to what children may show towards other interactive toys. Once children discover (through play and inquiry from the researchers) that KASPAR has a wider range of abilities than conventional interactive toys that can be bought in toy shops, their curiosity appears to get reinforced and they continue to engage with KASPAR more systematically, e.g. exploring its eyes etc. For typically developing children the minimal/subtle expressiveness in KASPAR seems to encourage them to reply

Figure 9. Some of KASPAR's expressions. Children usually interpret these expressions as 'good bye' (top, left), 'happy' (top, middle), 'surprised' (top, right), 'sad' (bottom, left) and 'thinking' (bottom, right). Note, our goal was not to create scientifically plausible emotional and other expressions (compare FEELIX, Kismet), but to create a robot with – from a user-centred perspective – appealing and interactional salient features.

with emphasised or bigger expressions in return, i.e. with a bigger smile, and bigger hand movements in imitation games etc.

- Adults show in general a more cautious and less playful attitude towards KASPAR, sometimes commenting on specific design features, e.g. noticing that the head is disproportionately larger than the rest of its body (as has been explained, this was a deliberate cartoon-inspired design choice). It appears (from explicit comments given to the researchers) that adults tend to spontaneously compare KASPAR with very realistically human-like robots they have seen in movies or on television. Their expectations towards the robot's capabilities are similarly high, so overall, adults tend to have a more critical attitude towards the robot. For these reasons, in our experiments involving adult participants we took care to introduce the robot and its capabilities before the start of the experiment.

Psychologists may further investigate the above issues, which go beyond the scope of our research, in future systematic studies.

4. Applications of KASPAR in research

Since 2005 our research team has been using KASPAR extensively in various research projects in the area of robot-assisted play, developmental robotics, gesture communication and development and learning. This section illustrates the experiments and the results that were obtained from some of these studies. We discuss these studies in the light of KASPAR's interaction abilities that afford a great variety of different human–robot interaction experiments. Note, a detailed description of the motivation, research questions, experiments and results would go beyond the scope of this paper. Instead, the following sections aim to *illustrate* the different usages of the robot in different interaction scenarios and applications where different methodological approaches have been used in the research and to document the experiments. Case study I illustrates work in a project in assistive technology based on case study evaluations whereby a narrative format has been chosen to describe the work. Case study II is situated in the context of human–robot interaction studies whereby a more experimental approach has been taken that takes into account not only the evaluation of the performance of the human–robot dyad (pair) but also the subjective evaluations of the experiment participants. Finally, case study III reports on research in developmental robotics whereby the emphasis is on the development and evaluation of cognitive architectures for robot development that relies on human interaction. Each case study will provide pointers to published work on these experiments so that the reader is able to find detailed information about the different methodological approaches, experiments and results.

4.1. Case study I: robot-assisted play and therapy

This case study discusses the use of KASPAR in robot-assisted play, in the specific application context of therapy for children with autism.

4.1.1. Motivation

Our research group has been involved for more than 10 years in studies that investigate the potential use of robots in autism therapy (Dautenhahn and Werry 2004) as part of the Aurora project (Aurora 2008). Different humanoid as well as non-humanoid robots have been used. The use of robots in robot-assisted play (with therapeutic and/or educational goals) is a very active area of research and a variety of special-purpose robots have been developed in this area (Michaud et al. 2003; Kozima et al. 2005; Saldien et al. 2008). Other work is exploring available research platforms (Kanda and Ishiguro 2005; Billard et al. 2006) or commercially available robots in an educational context (Tanaka et al. 2007). While in the area of assistive technology a variety of special requirements and needs need to be considered (cf. Robins et al. 2007 which reports on the IROMEC project that specifically designs a novel robot for the purpose of robot-assisted play for children who cannot play), KASPAR originally had not been designed only for this specific application area. However, as discussed above, the design of KASPAR included lessons learnt from the use of robots in autism therapy. And not surprisingly, KASPAR turned out to be a very engaging tool for children with autism and has been used extensively as an experimental platform in this area over the past few years.

This section presents some case study examples that highlight the use of KASPAR in the application area of autism therapy. Autism here refers to Autistic Spectrum Disorders, a range of manifestations of a disorder that can occur to different degrees in a variety of forms (Jordan 1999). The main impairments that are characteristic of people with autism, according to the National Autistic Society (NAS 2008), are impairments in social interaction, social communication and social imagination. This can manifest itself in difficulties in understanding gesture and facial expressions, difficulties in forming social relationships, the inability to understand others' intentions, feelings and mental states etc. They also usually show little reciprocal use of eye contact. As people's social behaviour can be very complex and subtle, for a person with deficits in mind-reading skills (as with autism), this social interaction can appear widely unpredictable and very difficult to understand and interpret.

KASPAR, which was designed as a minimally expressive humanoid robot, can address some of these difficulties by providing a simplified, safe, predictable and reliable environment. The robot was found to be very attractive to children with autism and a suitable tool to be used in education and therapy. As autism can manifest itself to different degrees and in a variety of forms, not only children in different schools might have different needs but also children in the same school might show completely different patterns of behaviour from one another and might have different or even some contradictory needs. Importantly, interaction with KASPAR provides multi-modal embodied interaction where the complexity of interaction can be controlled and tailored to the needs of the individual child and can be increased gradually.

4.1.2. Illustration of trials

The following examples show the potential use of KASPAR in education and therapy of children with autism. They present a varied range of settings (e.g. schools, therapy sessions etc.) and children who vary widely in their abilities and needs (from very low functioning children to high functioning and those with Asperger syndrome). KASPAR was found to be very attractive to all these children regardless of their ability. Children who were usually not able to tolerate playing with other children initially used KASPAR in solitary play and closely explored its behaviour, postures and facial features and expressions. Later, assuming the role of a social mediator (Robins et al. 2004b; Marti et al. 2005) and an object of shared attention, KASPAR helped these children (and others) in fostering basic social interaction skills (using turn-taking and imitation games), encouraging interaction with other children and adults. All trials took place in schools for children with special needs (Examples I–V) or health centres (Example VI). The experimenter was part of and actively involved in all of the trials; compare with Robins and Dautenhahn (2006) for a detailed discussion on the role of the experimenter in robot-assisted play.

The examples in school are part of a long-term study where children repeatedly interact with KASPAR over several months. More details about trials and analysis of the results can be found in Robins et al. (2009).

Example I. KASPAR promotes body awareness and sense of self

KASPAR encourages tactile exploration of its body by children of different age groups irrespective of their gender (Figure 10). All children with autism who first met KASPAR were drawn into exploring him in a very physical way. This tactile exploration is important to increase body awareness and sense of self in children with autism.

Example II. KASPAR evokes excitement, enjoyment and sharing – mediates child/adult interaction

We observed situations when children with severe autism who have very limited or no language at all got excited in their interaction with KASPAR and sought to share this experience with their teachers and therapists. These

Figure 10. Tactile exploration of KASPAR by children from different age groups and gender.

human contacts may give significance and meaning to the experiences with the robot (Figure 11).

Example III. KASPAR helps to break the isolation

Liam is a child with severe autism. Although in his home he interacts regularly with other family members, in school he is withdrawn to his own world, not initiating any interaction with other people (neither with other children nor with the teachers). After playing with KASPAR once a week for several weeks, Liam started to share his experience with his teacher (Figure 11, left), exploring the environment and communicating (in a non-verbal manner) with adults around him (both with the teacher and the experimenter) as can be seen in Figures 12 and 13.

Example IV. KASPAR helps children with autism to manage collaborative play

KASPAR's minimal expressiveness, simple operation and the use of a remote control encourage children not only to play with it but also to initiate, control and manage collaborative play with other children and adults (see Figures 14 and 15).

Example V. KASPAR as a tool in the hands of a therapist

As stated above, interaction with KASPAR is a multimodal embodied interaction where the complexity of interaction can be controlled, tailored and gradually increased to the needs of the individual child. Figure 16 shows how a therapist is using KASPAR to teach a child with severe autism turn-taking skills. Adam is a teenager who does not tolerate any other children, usually his focus and attention lasts only for very short time, he can be violent towards others and can also cause self-injury. However, after he was first introduced to KASPAR, he was completely relaxed, handled KASPAR very gently and kept his attention focused on it for as long as he was allowed (approximately 15 minutes). The therapist used his keen interest in KASPAR to teach him turn-taking skills with another person. Initially, Adam insisted on being in control all the time and refused to share KASPAR with anyone else, but after a while he allowed the therapist to take control, and slowly they progressed into full turn-taking and imitation games.

Example VI. KASPAR as a teaching tool for social skills

KASPAR was used in a pilot scheme to teach children with autism social skills during their family group therapy sessions run by the local child and adolescent mental health centre. During these sessions children practise how to approach other children to befriend them in the playground and in school. Children learnt how to ask precise questions by approaching KASPAR (as a mediator between them and other children), asking the robot a question and interpreting its response. KASPAR was operated by another child who gave the answer indirectly via the robot's gestures and facial expressions (Figure 17).

Figure 11. Liam seeks to share his excitement with his teacher (left); Derek shares his enjoyment with his therapists (right).

Figure 12. Liam is exploring KASPAR's facial features very closely (in this snapshot it concerns the eyes) and then turns to his teacher and explores her face in similar way.

Example VII. Use of a remote control by children with autism to operate KASPAR

In Examples IV and VI children used the remote control (Figure 18) to facilitate collaborative play. They were given the remote control and shown how to operate it. Most children got excited once they discovered and explored the use of its keypad, and asked for it every time they came to play with KASPAR.

The objectives for the children to use the remote control were varied. For those children who always wanted to be in control (a typical behaviour in autism), the remote control was a tool for learning turn-taking skills. It was a 'reward' once they learnt to 'let go' of the control, and not only gave it to another person but also participated in an imitation game where the other person was controlling the robot. For children who are usually passive and follow any instruction given, the use of the remote control encouraged taking initiative, discovering cause and effect and realising that they could also do actions on their own (e.g. they can change the robot's posture).

Moreover, whenever possible, the experimenter and a child, or two children were encouraged to play together (e.g. an imitation game), with the robot assuming the role of a social mediator. In this scenario the remote control is a key object that facilitates the acquisition of new skills that are *vital* for children with autism, i.e. they no longer merely follow instructions of games given to them by adults (which is often the case in classroom settings) but are also actually allowed *to take control of a collaborative game* to initiate, follow, take turns and even have the opportunity to give instructions to their peers.

Figure 13. Liam communicates with the experimenter.

Figure 14. Billy controls an imitation game (using a remote control) in a triadic interaction with the robot and the experimenter.

Figure 15. KASPAR mediates child–child interaction in a turn-taking and imitation game: one child controls KASPAR via remote control, the other imitates KASPAR. The children then switch roles.

4.1.3. *Reflections on KASPAR's design*

As has been mentioned, the Aurora research team has been using a variety of different robots in robot-assisted play for children with autism, including non-humanoid mobile robots, a humanoid robotic doll as well as a zoomorphic (in this case dog-like) robot (see Figure 19).

All three approaches with different robots used have in common that the child's control of the robot is indirect, i.e. through interaction – the robot and the child are active participants in the interaction, and enjoyment of the child is a key aim. Also, in all three studies the child can influence whatever game is being played. Table 1 shows in boldface the specific features of KASPAR that have turned out to be very successful during interactions with children with autism, as demonstrated in the above-mentioned case studies.

To summarise, following are the key features of KASPAR that turned out to be very important in the robot-assisted therapy with children with autism:

- A variety of facial/head and gestural expressions that allow a spectrum of social interaction and communicative as well as collaborative games.
- A remote control to operate the robot that can be operated by the experimenter or therapist as well as by children themselves. This control forms the basis of a variety of different games, e.g. imitation and turn-taking games.
- The remote control-facilitated collaborative games among children on their own initiative.

Note, after reviewing the literature (see discussion in Dautenhahn and Werry 2004) and discussions with psychologists we suggest that some of the attractiveness of KASPAR to children with autism is its minimal expressiveness, i.e. possessing simple facial features with less details – a face that appears less overwhelming and thus less threatening to children (in comparison to a person's face with numerous facial details and expressions that often are overwhelming to children with autism causing information overload). Also, KASPAR's limited amount of facial expressions makes its behaviours more predictable, which again suits the cognitive needs of children with autism. The generally very positive reactions from children (some verbal but most non-verbal due to limited language abilities) further support the view that KASPAR can provide a safe and enjoyable interactive learning environment for children with autism as motivated in Section 4.1.1.

4.2. *Case study II: drumming with KASPAR – studying human–humanoid gesture communication*

This second case study concerns the use of KASPAR in the European project 'Robotic Open-Architecture Technology for Cognition, Understanding, and Behaviours' (Robotcub; Sandini et al. 2004; Robotcub 2008) in the field of developmental robotics.

4.2.1. *Motivation*

'[I]nterpersonal coordination is present in nearly all aspects of our social lives, helping us to negotiate our daily face-to-face encounters . . . We also coordinate our nonverbal behavior with others to communicate that we are listening to them and want to hear more' (Bernieri and Rosenthal 1991, p. 401).

Figure 16. A therapist is using KASPAR to teach turn-taking skills to a child with autism.

Figure 17. KASPAR as part of family group therapy sessions to mediate between children and teach social skills.

Over the past two years KASPAR has been used extensively in our *drum-mate* studies, which investigate the playful interaction of people with KASPAR in the context of drumming games as a tool for the study of non-verbal communication (Kose-Bagci et al. 2007, 2008a, 2008b). This work forms part of our studies on gesture communication as part of the EU 6th framework project Robotcub. Drumming is a very suitable tool to study human–humanoid non-verbal communication because it includes issues such as social interaction, synchronisation, and turn-taking which are important in human–human interaction (Kendon 1970; Hall 1983; Bernieri and Rosenthal 1991; Goldin-Meadow and Wagner 2005). In robotics, different works have used robot drumming as a test bed for robot controllers (Kotosaka and Schaal 2001; Degallier et al. 2006). Other approaches focus on the development of a robot drummer that is able to play collaboratively with professional musicians (Weinberg et al. 2005; Weinberg and Driscoll 2007) or in concert with human drummers and at the direction of a human conductor (Crick et al. 2006). Our work uses drumming as a test bed for the study of human–humanoid non-verbal interaction and gesture communication.

From a practical viewpoint, drumming is relatively straightforward to implement and test, and can be applied technically without special actuators like fingers or special skills or abilities specific to drumming. So we could implement it with the current design of KASPAR, without additional need for fingers, or extra joints. With just the addition of external microphones for sound detection, it was able to perform drumming with tambourine style toy drums (Figure 20). Note, we did not need an additional drumstick, as due to its specific design KASPAR's hands are able to perform the drumming. In these experiments only one hand (the left one) was used for the drumming.

Figure 18. The remote control used in scenarios with children with autism.

4.2.2. *Drumming experiments with KASPAR*

KASPAR, in our experiments, has the role of an autonomous 'drumming companion' in call-and-response games, where its goal is to imitate the human partner's drumming (Figure 20). In the drum-mate studies, the human partner plays a rhythm, which KASPAR tries to replicate,

Figure 19. Top row: Non-humanoid, mobile robots used in the Aurora project – Aibo (left, Sony), Labo-1 (right, AAI Canada, Inc.). Bottom row: Different appearances of Robota, the humanoid doll–robot that has been used with children with autism. The 'robot-like' appearance on the right has been shown to be more engaging in first encounters of children with autism compared to Robota, the doll-robot (Robins et al. 2006).

Figure 20. A screen shot from the experiments where KASPAR is a drum-mate of human interaction partners.

in a simple form of imitation (mirroring[5]). KASPAR has two modes: listening and playing. In the listening mode, it records and analyses the played rhythm, and in the playing mode, it plays the rhythm back by hitting the drum positioned in its lap. Then the human partner plays again. This turn taking will continue for the fixed duration of the game. KASPAR does not imitate the strength of the beats but only the number of beats and duration between beats, due to its limited motor skills. It tailors the beats beyond its skills to those values allowed by its joints. KASPAR needs a small time duration (e.g. at least 0.3 seconds in the experiments) between each beat to get its joints 'ready', so that even if the human plays faster, KASPAR's imitations will be slower using durations of at least 0.3 seconds between beats. It also needs to wait for a few seconds before playing any rhythm in order to get its joints into correct reference positions.

In the first set of experiments (Kose-Bagci et al. 2007), head gestures accompanied the drumming of KASPAR. Here KASPAR just repeated the beats produced by the human partner, and made simple fixed head gestures accompanying its drumming (we used very simple gestures, without overt affective components like smiling or frowning in order not to overly distract the participants during the experiments). The participants, in return, perceived these simple behaviours as more complex and meaningful and adapted their behaviour to the robot's gestures. In this part of the study, we used deterministic turn-taking skills, simply mirroring the human's playing, which caused problems in terms of timing and negatively affected human participants' enjoyment. In the second part of the study (Kose-Bagci et al. 2008a), we developed novel turn-taking methods that appear more natural and engage the human participants more positively in the interaction games. Here, computational probabilistic models were used to regulate turn-taking skills of KASPAR emerging from the dynamics of social interaction between the robot and the human partner. Although we used very simple computational models, and this work is a first step in this domain, we were able to observe some very 'natural' games in terms of coordinated turn-taking games, and some of the participants even compared the game to a game they might play with children.

From the first set of experiments and our public demonstrations where we used gestures as social cues, we got positive feedback from the participants (48 adults and 68 primary school children). Especially at the public demonstrations where we used more complex gestures (e.g. smiles when KASPAR imitated human drumming, frowns when KASPAR could not detect human drumming or waving 'good bye' with a big frown when it had to finish the game), we got very positive feedback and public attention.

The reason behind KASPAR's successful head and face gestures is hidden in its face design. KASPAR's facial expressions and head and arm gestures seemed to influence the way human participants perceive the robot and the interaction. Even blinking and nodding and other head movements affect significantly human participants' evaluations of the robot and the games. Besides, the size of KASPAR makes it appear more 'child-like' which affects people's evaluations. Some of the adult participants compared the drumming experience they had with KASPAR with the experiences they had with their two to three-year-old children.

It is important to note that while KASPAR's drum playing did not change over time, and stayed the same in different games, the participants learned the limits of KASPAR and the rules of the game. Participants seemed to adapt themselves to the game better and the success rate improved over time. Humans, as shown here, were not passive subjects in this game, but adapted themselves to the capabilities of the robot. In order to facilitate and motivate such an adaptation, aspects of the interaction that are not directly related to the task itself, such as interactional gestures – like KASPAR's simple head/face gestures and blinking – may play an important role. A variety of research questions have been addressed using KASPAR in human–robot drumming experiments. A detailed discussion and results pertaining to these questions would go beyond the scope of this paper but can be found in Kose-Bagci et al. (2007, 2008a, 2008b). The next section illustrates some of the results.

4.2.3. Results and discussion

The following is a brief summary of results of some of the key points resulting from experiments presented in Kose-Bagci et al. (2007, 2008a, 2008b).

[5]Here we use 'mirroring' to refer to generalised matching of aspects of behaviour in interaction, e.g. number and timing of beats in a drumming interaction. In particular, it does not refer here to ipsilateral vs. contralateral imitation. Mirroring plays an important part in communicative interactions and the social development of children. For further discussion of mirroring and imitation, see Nehaniv and Dautenhahn (2007) and Nadel and Butterworth (1999).

Table 1. Design space of robots explored in the Aurora project: a comparison of three approaches with different robots. Also see related comparisons in Davis et al. (2005).

	Labo-1 (Werry and Dautenhahn 2007; Dautenhahn 2007)	**Robota** (Robins et al. 2004a, 2004b, 2006; Dautenhahn and Billard 2002)	**KASPAR** (see case studies)
Appearance	Mechanical-looking	Doll or plain appearance	Humanoid
Mode of operation	Autonomous	Remote-controlled	Remote-controlled
Mobility	Movements in 2-D on the floor (translational and rotational movements)	Movements of head (left, right), lifting of arms and legs (up, down)	**Different movements of the head/neck, different facial expressions (e.g. 'surprised', 'happy', 'sad' etc.), variety of arm gestures (e.g. waving, peek-a-boo etc.)**
Tasks with objects	Indirectly (obstacle avoidance)	None	Drumming (playing a toy tambourine)
Spatial dimensions of interaction	3-D, the child can approach and interact with the robot from any direction, child can also pick robot up etc.	3-D, but the child must be positioned in front of the robot to interact with it	3-D, but the child must be positioned in front of the robot to interact with it
Systems behaviours used	1. Few predetermined behaviours and simple action–selection architecture based on the robot's sensory input and internal states 2. Emergent, i.e. behaviours emerge from the interaction of the robot with the environment	Few predetermined behaviours elicited under control of a puppeteer who selects the robot's actions based on his perception of the situation and knowledge about the child, the interaction history/context etc.	Few predetermined behaviours elicited under control of a puppeteer who selects the robot's actions based on his perception of the situation and knowledge about the child, the interaction history/context etc.
Stance and movement of child during interaction with the robot	Child is free to run around the room, sit or crawl on the floor, approach, follow, avoid or pick up the robot	Child is free to sit, stand, move towards or away from the robot and touch it	Child is free to sit, stand, move towards or away from the robot and touch it
Control over the robot by child	Indirectly, through interaction	Indirectly, through interaction	1. Indirectly, through interaction **2. Child can use a remote control to operate the robot**
Nature of the interaction	Free, playful, unstructured, basic turn-taking and approach/avoidance routines lead to games such as following, chasing etc.	Free interaction, but guided by experimenter, e.g. 'look at what the robot/the other child is doing'	1. Free interaction, but guided by experimenter, e.g. 'look at what the robot/the other child is doing' **2. By controlling the robot via a remote control, the child can *manage* a collaborative game with another child on his/her own initiative**
Targeted therapeutic behaviours	Turn-taking, joint attention, proactive behaviour, initiative taking, mediation between child and other persons via the robot	Turn-taking, joint attention, imitation of limb movements, proactive behaviour, initiative taking, mediation between child and other persons via the robot	Turn-taking, joint attention, collaborative activities, imitation of hand gestures and head and facial expressions, proactive behaviour, initiative taking, mediation between child and other persons via the robot, body awareness & sense of self
Tailoring to needs of individual children	No individual adaptation was used	Manual adaptation by puppeteering	Manual adaptation by puppeteering

- A trade-off between the subjective evaluation of the drumming experience from the perspective of the participants and the objective evaluation of the drumming performance. Participants preferred a certain amount of robot gestures as a motivating factor in the drumming games that provided an experience of social interaction. However, the sample was divided in terms of what degrees of gestures were appropriate.
- The more games participants played with the robot the more familiar they became with the robot; however, boredom was also mentioned by some participants which indicates the essential role of research into how to maintain a user's interest in the interaction with a robot.
- The more participants played with the robot the more they synchronised their own drumming behaviour with the robot's. The different probabilistic models that controlled the robot's interaction dynamics led to different subjective evaluations of the participants and different performances of the games. Participants preferred models that enable the robot and humans to interact more and provide turn-taking skills closer to 'natural' human–human conversations, despite differences in objective measures of drumming behaviour. Overall, results from our studies are consistent with the *temporal behaviour-matching hypothesis* previously proposed in the literature (Robins et al. 2008), which concerns the effect that participants adapt their own interaction dynamics to that of the robot's.

4.2.4. *Reflections on KASPAR's design*

How suitable has been KASPAR in the interaction experiments using drumming games? KASPAR's movements do not have the precision or speed of industrial robots or some other humanoid robots that have been developed specifically for manipulation etc. One example of a high-specification robot is the iCub that has been developed within the European project Robotcub at a cost of €200,000 (Figure 21). The iCub has the size of a 3.5-year-old child, is 104 cm tall and weighs 22 kg. It has 53 joints mainly distributed in the upper part of the body. While KASPAR has been designed from off-the-shelf components, every component of the iCub has been specifically designed or customised for the robot in order to represent cutting edge robotic technology.

Also, special purpose robotic percussionists have been designed specifically for the purpose of *efficient* drumming, e.g. Haile (Weinberg et al. 2005). The design rationale of Haile, a robot with an anthropomorphic, yet abstract shape that can achieve drumming speeds of up to 15 Hz, was very different from KASPAR: 'The design was purely functional and did not communicate the idea that it could interact with humans by listening, analyzing, and reacting' (Weinberg and Driscoll 2007). Haile is a special purpose drumming robot that can join and improvise with live professional players. Unlike Haile, which was specifically designed for performing drumming, KASPAR is using drumming as a *tool for social interaction*. Detailed technical comparisons of KASPAR with Haile or the iCub are not useful because these all serve very different purposes. For example, the iCub has been designed for tasks such as crawling and manipulation, and Haile can achieve impressive drumming performances in terms of speed and precision.

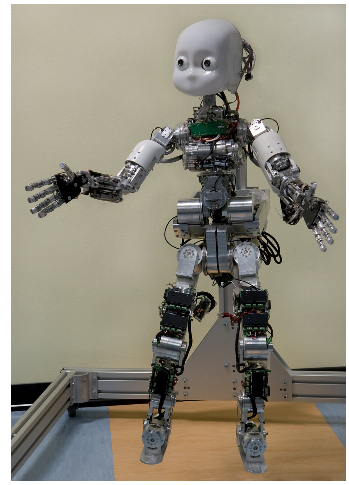

Figure 21. The iCub robot.

However, despite KASPAR's low-precision design, our studies have shown that it is very suitable for human–robot interaction studies where speed, precision or complex movement patterns are not of primary importance, as is the case in our experiments on drumming games that were successful in terms of social interaction, imitation and turn-taking. And it is in such cases that the low-cost robot KASPAR, which can easily be built and maintained by robotics researchers, is *socially* effective and suitable as a tool for interaction experiments. Also, compared with the iCub, KASPAR is safer to use in interactions even when involving children and tactile interactions with people (cf. Section 4.1.2 where, in the case of children with autism interacting with KASPAR they often touched the robot, e.g. stroking or squeezing the cheeks, tapping the chin etc.). KASPAR moves relatively slowly and cannot exhibit strong forces, which limits the risks involved in human–robot interaction[6]. Even small children can easily stop, e.g. KASPAR's arm movements by simply grabbing its hands/arms, and the coverage of metal parts

[6] We believe that any device or toy used in interactions with people can potentially provide a safety risk, e.g. children can choke on CE-certified commercially available toys. Thus, it is a matter of reducing risks as much as possible.

with clothing (or parts of the original mannequin used for the hands) prevents cuts and bruises.

4.3. *Case study III: 'peekaboo' – studying cognition and learning with KASPAR*

This last case study illustrates the use of KASPAR II, as part of the above-mentioned project Robotcub, for the investigation of cognition and learning. In this section we provide a brief summary of this research illustrating the use of KASPAR II. More details about this particular experiment can be found in Mirza et al. (2008).

4.3.1. *Motivation*

Why use a robot to study cognition? The answer to this question defines modern research into Artificial Intelligence and the mechanisms and processes that contribute to the cognitive capabilities of humans and many other animals. Increasingly, the importance of embodiment and situatedness within complex and rich environments are becoming recognised as crucially important factors in engendering intelligence in an artifact (see for example Clancey 1997; Pfeifer and Bongard 2007) and the philosophical position regarding 'structural coupling' of Maturana and Varela (1987). The 'embodied cognition' hypothesis argues that 'cognition is a highly embodied or situated activity and suggests that thinking beings ought therefore be considered first and foremost as acting beings' (Anderson 2003).

That many aspects of cognition are grounded in embodiment is not the whole story though. We want to take a further step and ask 'why use a humanoid robot with expressive capabilities to study cognition?' In this case, two other aspects come into play. Firstly, having a human-like body allows the robot to participate in a social context, and secondly, in the absence of a language, being able to evoke emotional responses in a human interaction partner through facial expressions, the communicative capability of the robot is greatly enhanced.

In this section we describe research work that uses the early-communicative interaction game 'peekaboo' as a scenario through which aspects of ontogenetic development (i.e., development over a lifetime through accumulation of experience) can be studied. The research is focused on understanding how an *interaction history* (Mirza et al. 2007), developed continually over time from the sensorimotor experience of a robot, can be used in the selection of actions in playing the 'peekaboo' game.

'Peekaboo' is a well-known interaction game between infant and caregiver where, classically, the caregiver, having established mutual engagement through eye contact, hides their face momentarily. On revealing their face again the caregiver cries 'peekaboo!', or something similar usually resulting in pleasure for the infant and cyclic continuation of the game. Bruner and Sherwood (1975) studied the game in terms of its communicative aspects showing that timing is crucial. Moreover, research shows that such games can serve as scaffolding for the development of primary intersubjectivity and the co-regulation of emotional expressions with others (Rochat et al. 1999).

4.3.2. *'Peekaboo' experiments with KASPAR*

In order to better understand the experiments, we first provide brief details on the robot's interaction history architecture and its socially interactive behaviour. More information about the experiments and results are provided in Mirza et al. (2008).

4.3.2.1. Interaction history architecture. The interaction history architecture has at its heart a mechanism for relating the continuous sensorimotor experiences of a robot in terms of their information-theoretic similarity to one another. At any time the robot's current experience (in terms of the sum of its sensorimotor values for a given period of time, the time-horizon h) can be compared to those in its history of interaction. The most similar one from the past can then be used to extract an action policy that was earlier successful. The feedback from the environment acts to enhance those experiences that result in high reward for the robot. By bootstrapping the history, by exploring interaction possibilities, by executing any action from its repertoire, the robot can rapidly develop the capability to act appropriately in a given situation. See Mirza et al. (2007) for further details.

4.3.2.2. Actions, feedback and reward. A total of 17 actions were available to the robot, and these can be considered in three groups: movement actions (e.g head-right, wave-right-arm or hide-head), facial expressions (e.g. smile – see Figure 22) and resetting actions (e.g. reset)[7]. The facially expressive actions convey the response of the robot in terms of the reward it receives. This provides instantaneous feedback for the interaction partner. Reward is given as an integral part of the interaction. The human partner encourages the robot with calls of 'peekaboo'. Such an increase in sound level combined with the detection of a face by the robot's camera-eyes, results in a high reward.

4.3.2.3. Experimental method. The robot faces the human partner and the interaction history started, initially empty of any previous experience. Interaction then commences with the robot executing various actions and the human offering vocal encouragement when thought appropriate, which continues for about three minutes. Three

[7]The actions that can be executed at any time are restricted for reasons of practical safety of the robot.

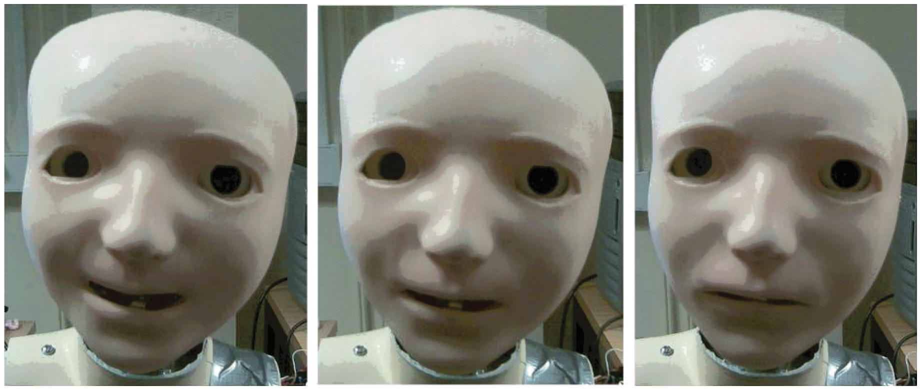

Figure 22. Facial expressions of KASPAR II. From left –to right: smile, neutral, frown.

different conditions were tried. Firstly, the hid-face behaviour was encouraged with a call of ‘peekaboo’ when the robot revealed its face again. The second condition encouraged an alternative action (such as turn-head-left) and the final condition was to offer no vocal encouragement at all during the interaction.

4.3.3. Results and discussion

A total of 22 runs were completed. Sixteen of these for the first condition (encouraging the hiding action), three for the second one and three for the no-encouragement condition. In 67% of the cases where reward was given (‘peekaboo’ or otherwise), the robot repeated the encouraged behaviour. In the cases where no encouragement was given no repeated action took place.

Figure 23 shows for the first run (d0032), how the motivational variables (face, sound and resultant reward) vary with time, along with the actions being executed. The interaction partner encourages the first ‘peekaboo’ sequence (‘hide-face’ on the diagram). Note that the ‘peekaboo’ behaviour is actually a combination of actions to hide

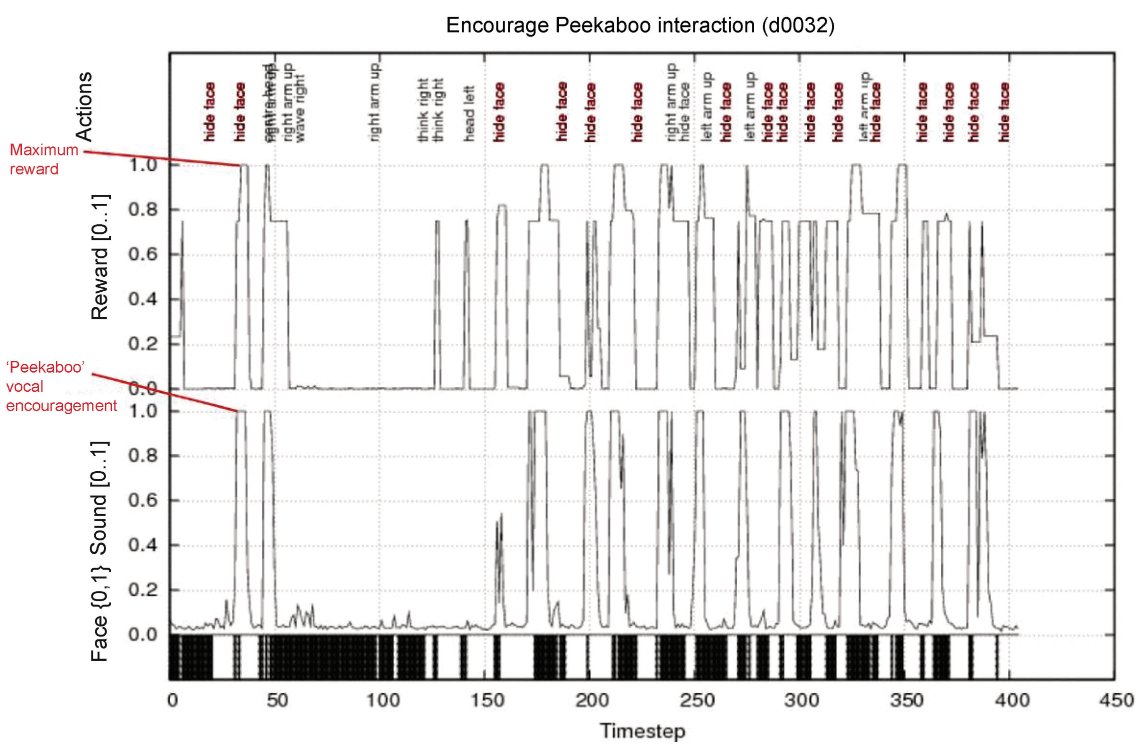

Figure 23. Illustration of results: Example of ‘peekaboo’ encouragement condition. The trace shows, against time, the detection of the face and audio encouragement as well as the resulting reward. Along the top are shown the actions executed.

the face (action 6), any number of 'no-action' actions (action 7) and an action to return to the forward resting position (action 0) (for clarity only the primary action is shown on the trace). This results in a maximal reward shortly after the hide-face action, and as the interaction partner continues to reinforce the 'peekaboo' behaviour with vocal reward, this pattern can be seen repeated throughout the trace.

The results supported the hypothesis that by encouraging the behaviour the interaction history of the robot would cause combinations of actions to be repeated in search of more reward. Furthermore, the exact combination of actions necessary is not hard coded as other action combinations can be similarly encouraged. Finally, not providing encouragement results in random, non-interactive behaviour. It was also found that the timings of the feedback and thus the interaction were important – too early or too late and alternative actions were encouraged.

4.3.4. Reflections on KASPAR's design

Any embodied agent engaging in temporally extended interaction with its environment can make use of an interaction history; however, the particular embodiment plays an important role in managing both the types of interactions that are possible as well as the expectations of such possibilities in an interaction partner. As such, the particular design of the KASPAR series of robots plays an important role. For instance, bearing a physical similarity to that of a human infant means that complex speech will not be expected, but that attention to a human face and sounds might be expected. Probably, the most important aspect of the physical design of KASPAR is its expressive face that provides a mechanism for the robot's actuators to influence a human interaction partner just as a robotic arm might influence the position of an object. However, in terms of the interaction history, it is also important that the embodiment provides not only suitable actuators and appearance but also well-engaged sensory surfaces. These are crucial for providing information about how the environment is changing with respect to the actions of the robot. As such, the KASPAR robots provide both visual and auditory sensors as well as (in KASPAR II) proprioceptive sensors that feed back information about the positions of its joints over and above the controlled position. Overall, this experiment illustrated the suitability of the robot for quantitative experiments in cognitive and developmental robotics for research involving human–robot interaction scenarios where accuracy and speed of movements is not of primary importance.

5. Conclusion

This paper has described the development of a minimally expressive humanoid robot – KASPAR. The design rationale, guidelines and requirements, as well as the design of the robot itself were described in detail. We also discussed our approach in the context of related research work on socially interactive robots. While a detailed comparison of KASPAR with other robots, as well as experimental investigation comparing the suitability of those robots in human–robot interaction studies, go beyond the scope of this paper, in the following part we conceptually assess KASPAR (see Figure 24) according to different continuous scales ranging from high to low. We propose these dimensions as relevant assessment criteria for the design of humanoid (or other) robots used for multiple purposes involving interaction with people.

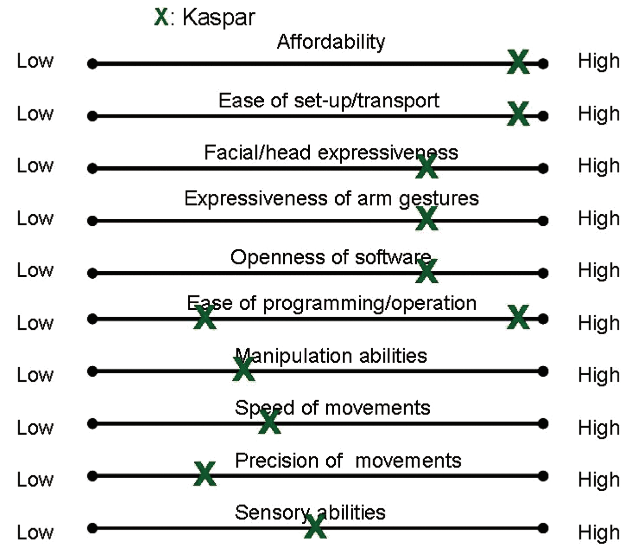

Figure 24. Assessment of the minimally expressive robot KASPAR. The continuous scales ranging from low to high provide a conceptual (not quantitative) assessment. Please note, for the 'ease of programming' category two estimations can be made, depending on whether one chooses to operate the robot in remote controlled mode/using the keypad (very easy to operate even by children), or whether the robot is used by researchers to develop new software (requires computer science knowledge).

KASPAR affords a variety of usages for human–robot interaction studies in the laboratory or in schools, in being able to provide a high degree of expressiveness and ability to carry out interaction games. Disadvantages of KASPAR concern the technical constraints on its movements in terms of speed, precision etc.; however, these issues are usually not crucial in more socially oriented human–robot interaction research. Note that the 'mobility' of KASPAR (i.e. ease of transport) and suitability for a variety of interaction scenarios (see section 4) and application areas are important to the field of human–robot interaction, as most existing robotic platforms are still limited to usage in the laboratory and need to be set up and operated by highly trained staff. KASPAR belongs to a new category of more 'user friendly' and (relatively) inexpensive robots that can be constructed by robotic students and researchers with no specific expert knowledge in humanoid robotics.

Generally, any robot designed for human–robot interaction scenarios is likely to have strengths and weaknesses depending on the particular requirements given by their application context. However, the assessment criteria proposed here may also be applicable to other robotic platforms and thus allow a matching of requirements posed by application contexts and robot abilities.

We hope that this paper has served multiple purposes:

- A detailed account of the design of a minimally expressive humanoid research platform that will inform other researchers interested in such designs.
- An introduction of key issues relevant in the design of socially interactive robots.
- An illustration of the use of the robot KASPAR in a variety of research projects ranging from basic research to application-oriented research.
- A discussion of the advantages and disadvantages of the socially interactive robot.

To conclude, designing socially interactive robots remains a challenging task. Depending on its envisaged purpose(s) different designs will be of different utility. Building a robot for a particular niche application is difficult, building a multi-purpose robot primarily for social interaction, as we did, is a huge challenge. The solution we found in KASPAR (and its offspring that already exists and new versions that are in the making) cannot be ideal, but it has served not only its original purpose but also exceeded our expectations to an unforeseen degree. The project to build KASPAR started in 2005 and was envisaged as a two-month, short-term project for a small study on humanoid expressiveness, and it was also the first attempt of our interdisciplinary research group to build a humanoid robot. We succeeded, as evidenced by a large number of peer-reviewed publications emerging from work with the robot. And KASPAR has been travelling the world to various conferences, exhibitions and therapy centres. But how to develop believable, socially interactive robots, in particular robots that can positively contribute to society as companions and assistants, remains a challenging (research) issue. We are still learning, and by writing this paper we would like to share our experiences with our peers.

Acknowledgements

We would like to thank Andrew Appleby who built the first prototype of KASPAR's head, as part of a design team and under the advice of Kerstin Dautenhahn, Chrystopher L. Nehaniv and David Lee. Michael L. Walters has been in charge of later hardware developments, with contributions from Mike Blow regarding the design of the arms. The keypad interface and the expressions shown in Figure 8 were designed by Ben Robins. Many other members of the Adaptive Systems research group, including authors of this paper, contributed to this work in a variety of ways including the software development of the robot. The research described in Sections 4.2 and 4.3 of this paper were conducted within the EU Integrated Project RobotCub and was funded by the European Commission through the E5 Unit (Cognition) of FP6-IST under Contract FP6-004370.

References

Anderson ML. 2003. Embodied cognition: a field guide. Artif Intell. 149:91–130.

Aryananda L. 2004. Mertz: a quest for a robust and scalable active vision humanoid head robot. Proceedings of IEEE-RAS International Conference on Humanoid Robots; 2004 Nov 10–12; Santa Monica, Los Angeles, CA. p. 513–532.

Aurora. 2008. Available from: www.aurora-project.com [cited 2008 Oct 15].

Bernieri FJ, Rosenthal R. 1991. Interpersonal coordination: behaviour matching and interactional synchrony. In: Feldman RS, Rimé B, editors. Fundamentals of nonverbal behavior. Cambridge: Cambridge University Press. p. 401–432.

Billard A, Robins B, Dautenhahn K, Nadel J. 2006. Building robota, a mini-humanoid robot for the rehabilitation of children with autism. RESNA Assistive Technol J. 19(1):37–49.

Blow MP, Dautenhahn K, Appleby A, Nehaniv CL, Lee D. 2006. Perception of robot smiles and dimensions for human-robot interaction design. Proceedings of the 15th IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN06); Sep 6–8; Hatfield, UK: University of Hertfordshire. IEEE Press. p. 469–474.

Breazeal CL. 2002. Designing sociable robots. Cambridge, MA: MIT Press.

Breazeal CL, Foerst A. 1999. Schmoozing with robots: exploring the boundary of the original wireless network. In: Proceedings of Third International Cognitive Technology Conference (CT-99); 1999; San Francisco, M.I.N.D. Lab, Michigan State University. p. 375–390.

Bruner JS, Sherwood V. 1975. Peekaboo and the learning of rule structures. In: Bruner J, Jolly A, Sylva K, editors. Play: its role in development and evolution. New York: Penguin. p. 277–285.

Cañamero LD. 2002. Playing the emotion game with feelix: what can a LEGO robot tell us about emotion? In: Dautenhahn K, Bond A, Cañamero LD, Edmonds B, editors. Socially intelligent agents: creating relationships with computers and robots. Norwell, MA: Kluwer Academic Publishers. p. 69–76.

Clancey WJ. 1997. Situated cognition: on human knowledge and computer representations. New York: Cambridge University Press.

Crick C, Munz M, Scassellati B. 2006. Synchronization in social tasks: robotic drumming. Proceedings of IEEE RO-MAN; 2006 Sep 6–8; Hatfield, UK: de Havilland Campus. p. 97–102.

Dautenhahn K. 1997. I could be you—the phenomenological dimension of social understanding. Cybernet Syst. 28(5):417–453.

Dautenhahn K. 1999. Embodiment and interaction in socially intelligent life-like agents. In: Nehaniv CL, editor. Computation for metaphors, analogy and agents. Springer lecture notes in computer science, vol. 1562. Berlin, Heidelberg, Germany: Springer Verlag. p. 102–142.

Dautenhahn K. 2002. Design spaces and niche spaces of believable social robots. In Proceedings of IEEE Intl. Workshop Robot and Human Interactive Communication (RO-MAN 2002); 2002 Sep 25–27; Berlin, Germany. p. 192–197.

Dautenhahn K. 2004. Socially intelligent agents in human primate culture. In: Payr S, Trappl R, editors. Agent culture:

human-agent interaction in a multicultural world. Philadelphia, PA: Lawrence Erlbaum Associates. p. 45–71.

Dautenhahn K. 2007. Socially intelligent robots: dimensions of human–robot interaction. Phil Trans R Soc B. 362(1480):679–704.

Dautenhahn K, Billard A. 2002. Games children with autism can play with Robota, a humanoid robotic doll, in universal access and assistive technology. In: Keates S. et al., editor. Proceedings of 1st Cambridge Workshop on Universal Access and Assistive Technology (incorporating 4th Cambridge Workshop on Rehabilitation Robotics); 2002 Mar 25–27; Trinity Hall, University of Cambridge, Cambridge, UK. London: Springer-Verlag. p. 179–190.

Dautenhahn K, Nehaniv CL. 2000. Living with socially intelligent agents: a cognitive technology view. In: Dautenhahn K, editor. Human cognition and social agent technology (advances in consciousness research series). Amsterdam, the Netherlands: John Benjamins Publishing. p. 415–426.

Dautenhahn K, Werry I. 2004. Towards interactive robots in autism therapy: background, motivation and challenges. Prag Cognition. 12(1):1–35.

Dautenhahn K, Woods S, Kaouri C, Walters M, Koay KL, Werry I. 2005. What is a robot companion—friend, assistant or butler? Proceedings of IROS 2005, IEEE IRS/RSJ International Conference on Intelligent Robots and Systems; 2005 Aug 2–6; Edmonton, Alberta Canada. p. 1488–1493.

Davis M, Robins B, Dautenhahn K, Nehaniv C, Powell SA. 2005. Comparison of interactive and robotic systems in therapy and education for children with autism. In: Assistive technology: from virtuality to reality—AAATE 2005. Lille, France. IOS Press. p. 353–357.

Degallier S, Santos CP, Righetti L, Ijspeert A. 2006. Movement generation using dynamical systems: a humanoid robot performing a drumming task. In Proceedings of the IEEE-RAS International Conference on Humanoid Robots (HUMANOIDS06); 2006 Dec 4–6; Genova, Italy.

DiSalvo C, Gemperle F, Forlizzi J, Kiesler S. 2002. All robots are not created equal: the design and perception of humanoid robot heads. In Proceedings of the 4th Conference on Designing Interactive Systems: Processes, Practices, Methods, And Techniques; London, UK. p. 321–326.

Ekman P, Davidson RJ, Friesen WV. 1990. The Duchenne smile: emotional expression and brain physiology II. J Pers Soc Psychol. 58(2):342–353.

Fong T, Nourbakhsh I, Dautenhahn K. 2003. A survey of socially interactive robots. Robot Auton Syst. 42(3–4):143–166.

Goldin-Meadow S, Wagner MS. 2005. How our hands help us learn. Trends Cogn Sci. 9(5):234–241.

Gong L, Nass C. 2007. When a talking-face computer agent is half-human and half-humanoid: human identity and consistency preference. J Hum Commun Res. 33(2):163–193.

Goodrich MA, Schultz AC. 2008. Human-robot interaction: a survey. Found Trends Hum-Comput Interact. 1(3):203–275.

Hall ET. 1983. The dance of life: the other dimension of time. New York: Doubleday/Anchor Press.

Johnson MJ, Morton J. 1991. Biology and cognitive development: the case of face recognition. Oxford, England: Blackwell.

Jones BC, Little AC, Burt DM, Perrett DI. 2004. When facial attractiveness is only skin deep. Perception. 33(5):569–576.

Jordan R. 1999. Autistic spectrum disorders—an introductory handbook for practitioners. London: David Fulton Publishers.

Kanda T, Ishiguro H. 2005. Communication robots for elementary schools. In: Proceedings of AISB'05 Symposium Robot Companions: Hard Problems and Open Challenges in Robot-Human Interaction; 2005 Apr; Hatfield Hertfordshire. p. 54–63.

Kendon A. 1970. Movement coordination in social interaction: some examples described. Acta Psychol. 32:100–125.

Kismet. 2009. Available from: http://www.ai.mit.edu/projects/humanoid-robotics-group/kismet/kismet.html [cited 2009 May 28].

Kose-Bagci H, Dautenhahn K, Nehaniv CL. 2008a. Emergent dynamics of turn-taking interaction in drumming games with a humanoid robot. Proceedings of IEEE RO-MAN 2008; 2008 Aug 1–3; Munich, Germany: Technical University Munich.

Kose-Bagci H, Dautenhahn K, Nehaniv CL. 2008b. Drumming with a humanoid robot: results from human-robot interaction studies. Proceedings of ECSIS Symposium on Learning and Adaptive Behavior in Robotic Systems, LAB-RS 2008; 2008 Aug 6–8 (published extended abstract of invited lecture); Edinburgh, Scotland. Washington, DC: IEEE Computer Society Press.

Kose-Bagci H, Dautenhahn K, Syrdal DS, and Nehaniv CL. 2007. Drum-mate: a human-humanoid drumming experience. Proceedings of IEEE-RAS Humanoids 2007; 2007 Nov 29–Dec 1; Pittsburgh, Pennsylvania, USA.

Kotosaka S, Schaal S. 2001. Synchronized robot drumming by neural oscillator. J Robot Soc Jpn. 19(1):116–123.

Kozima H, Nakagawa C, Yasuda Y. 2005. Designing and observing human-robot interactions for the study of social development and its disorders. Proceedings of 6th IEEE International Symposium on Computational Intelligence in Robotics and Automation—CIRA 2005; 2005 June 27–30; Espoo, Finland. p. 41–46.

Langlois J, Kalakanis L, Rubenstein A, Larson A, Hallam M, Smoot M. 2000. Maxims or myths of beauty? A meta-analytic and theoretical review. Psychol Bull. 126:390–423.

Langlois J, Roggman L. 1990. Attractive faces are only average. Psychol Sci. 1:115–121.

Lyons M, Campbell R, Plante A, Coleman M, Kamachi M, Akamatsu S. 2000. The Noh mask effect: vertical viewpoint dependence of facial expression perception. Proc R Soc Lond B. 267:2239–2245.

LynxMotion 2007. http://www.lynxmotion.com/Product.aspx?productID=395&CategoryID=52. Last accessed 14 July 2009.

MacDorman KF. 2005a. Androids as an experimental apparatus: why is there an uncanny valley and can we exploit it? In: CogSci-2005 workshop: toward social mechanisms of android science; Stresa, Piedmont, Italy. p. 106–118.

MacDorman KF, Ishiguro H. 2006. The uncanny advantage of using androids in social and cognitive science research. Interact Stud. 7(3):297–337.

MacDorman KF, Minato T, Shimada M, Itakura S, Cowley S, Ishiguro H. 2005b. Assessing human likeness by eye contact in an android testbed. In: Proceedings of 20th Annual Meeting of the Cognitive Science Society; 2005 July; Stresa, Italy.

MacDorman KF, Green RD, Ho C-C, Koch C. 2009. Too real for comfort: Uncanny responses to computer generated faces. Computers in Human Behaviour. 25:695–710.

Maturana HR, Varela FJ. 1987. The tree of knowledge: the biological roots of human understanding. Boston, MA: New Science Library (Shambhala).

McCloud S. 1993. Understanding comics: the invisible art. New York, NY: Harper Collins Publishers, Inc.

Michaud F, Duquette A, Nadeau I. 2003. Characteristics of mobile robotics toys for children with pervasive developmental disorders. IEEE Int Conf Syst Man Cybernet. 3:2938–2943.

Mirza NA, Nehaniv CL, Dautenhahn K, te Boekhorst R. 2007. Grounded sensorimotor interaction histories in an information theoretic metric space for robot ontogeny. Adapt Behav. 15(2):167–187.

Mirza NA, Nehaniv CL, Dautenhahn K, te Boekhorst R. 2008. Developing social action capabilities in a humanoid robot using an interaction history architecture. Proceedings of IEEE Humanoids 2008; 2008 Dec 1–3; Daejeon, Korea.

Miyake Y. 2003. Co-creation in human–computer interaction. In: Stephanidis C, Jacko J, editors, Human-computer interaction. London: Lawrence Erlbaum Associates Publishers. p. 513–517.

Mori M. 1970. Bukimi no tani [the uncanny valley]. Energy. 7:33–35.

Nadel J, Butterworth G, editors. 1999. Imitation in infancy. Cambridge: Cambridge University Press.

NAS. 2008. National Autistic Society, UK. Available from: http://www.nas.org.uk [cited 2008 Oct 9].

Nehaniv CL, Dautenhahn K, editors. 2007. Imitation and social learning in robots, humans and animals: behavioural, social and communicative dimensions. Cambridge: Cambridge University Press.

Pascalis O, Scott LS, Kelly DJ, Shannon RW, Nicholson E, Coleman M, Nelson CA. 2005. Plasticity of face processing in infancy. PNAS. 102(14):5297–5300.

Perrett D, May K, Yoshikawa S. 1994. Attractive characteristics of female faces: preference for non-average shape. Nature. 368:239–242.

Pfeifer R, Bongard J. 2007. How the body shapes the way we think: a new view of intelligence. Cambridge, MA: MIT Press.

Preece J, Rogers Y, Sharp H. 2002. Interaction design. New York, NY, USA: John Wiley & Sons, Inc.

Robins B, Dautenhahn K. 2006. The role of the experimenter in HRI research—a case study evaluation of children with autism interacting with a robotic toy. Proceedings of 15th IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN06); Sep 6–8; Hatfield, UK: University of Hertfordshire. IEEE Press. p. 646–651.

Robins B, Dautenhahn K, Dickerson P. 2009. From isolation to communication: a case study evaluation of robot assisted play for children with autism with a minimally expressive humanoid robot. Proceedings of Second International Conference on Advances in Computer-Human Interactions; ACHI 09; 2009 Feb 1–7; Cancun, Mexico. IEEE Computer Society Press. p. 205–211.

Robins B, Dautenhahn K, Dubowski J. 2006. Does appearance matter in the interaction of children with autism with a humanoid robot? Interact Stud. 7(3):509–542.

Robins B, Dautenhahn K, Nehaniv CL, Mirza NA, François D, Olsson L. 2005a. Sustaining interaction dynamics and engagement in dyadic child-robot interaction kinesics: lessons learnt from an exploratory study. Proceedings of 14th IEEE International Workshop on Robot and Human Interactive Communication (Ro-Man 2005); Aug 13–15; Nashville, USA. IEEE Press. p. 716–722.

Robins B, Dautenhahn K, te Boekhorst R, Billard A. 2004a. Effects of repeated exposure to a humanoid robot on children with autism. In: Proceedings of Universal Access and Assistive Technology (CWUAAT); 2004 Mar 22–24; Cambridge, UK. p. 225–236.

Robins B, Dautenhahn K, te Boekhorst R, Billard A. 2005b. Robotic assistants in therapy and education of children with autism: can a small humanoid robot help encourage social interaction skills? Special issue "design for a more inclusive world." Intern J Univers Access Inf Soc. (UAIS). 4(2):105–120. (Springer-Verlag).

Robins B, Dautenhahn K, te Boekhorst R, Nehaniv CL. 2008. Behaviour delay and robot expressiveness in child-robot interactions: a user study on interaction kinesics. Proceedings of ACM/IEEE 3rd International Conference on Human-Robot Interaction (HRI '08); New York, NY. ACM. p.17–24.

Robins B, Dickerson P, Stribling P, Dautenhahn K. 2004b. Robot-mediated joint attention in children with autism: a case study in robot-human interaction. Interact Stud. 5(2):161–198.

Robins B, Otero N, Ferrari E, Dautenhahn K. 2007. Eliciting requirements for a robotic toy for children with autism—results from user panels. Proceedings of 16th International Symposium on Robot and Human Interactive Communication (IEEE RO-MAN 2007); 2007 Aug 26–29; Jeju Island, Korea. p. 101–106.

Robotcub. 2008. Available from: http://www.robotcub.org [cited 2008 Oct 15].

Rochat P, Querido JG, Striano T. 1999. Emerging sensitivity to the timing and structure of protoconversation in early infancy. Dev Psychol. 35(4):950–957.

Saldien J, Goris K, Yilmazyildiz S, Verhelst W, Lefeber D. 2008. On the design of the huggable robot probo. Special issue on human interaction with domestic robots. J Phys Agents. 2(2):3–11.

Sandini G, Metta G, Vernon D. 2004. Robotcub: an open framework for research in embodied cognition. In: Proceedings of IEEE-RAS/RSJ Intl. Conf. Humanoid Robots (Humanoids 2004); 2004 Nov 10–12; Santa Monica, Los Angeles, CA.

Tanaka F, Cicourel A, Movellan JR. 2007. Socialization between toddlers and robots at an early childhood education center. Proc Natl Acad Sci USA. 104(46):17954–17958.

The Guardian Weekend 2007. Ophelia of the Seine. Saturday 1 December 2007. Available from: http://www.guardian.co.uk/world/2007/dec/01/france.art Last accessed 13 July 2009.

Wakamaru. 2009. Available from: http://www.mhi.co.jp/kobe/wakamaru/english/know/design/index.html [cited 2009 May 26].

Walters ML. 2008a. The design space for robot appearance and behaviour for social robot companions. PhD thesis. School of Computer Science, University of Hertfordshire. Available from: http://hdl.handle.net/2299/1806 [cited 2008 Oct 14].

Walters ML, Syrdal DS, Dautenhahn K, te Boekhorst R, Koay KL. 2008b. Avoiding the uncanny valley: robot appearance, personality and consistency of behavior in an attention-seeking home scenario for a robot companion. Auton Robots. 24(2):159–178.

Weinberg G, Driscoll S. 2007. The interactive robotic percussionist: new developments in form, mechanics, perception and interaction design. Proceedings of ACM/IEEE HR'I07; 2007 Mar 10–12; Arlington, Virginia. p. 97–104.

Weinberg G, Driscoll S, Parry M. 2005. Musical interactions with a perceptual robotic percussionist. Proceedings of IEEE International Workshop on Robot and Human Interactive Communication (RO-MAN 2005); Nashville, TN, USA. pp. 456–461.

Werry I, Dautenhahn K. 2007. Human-robot interaction as a model for autism therapy: an experimental study with children with autism. In: Laubichler M, Müller GB, editors. Modeling biology: structures, behaviors, evolution. Vienna series in theoretical biology. Cambridge, MA: MIT Press. p. 283–299.

Yarp. 2008. Available from: http://eris.liralab.it/yarp/ [cited 2008 Oct 15].

Appendix A

Head design and construction

The head was designed to mount and support the face mask and provide actuation for the facial expressions. The neck has three main DoFs: pan, tilt and roll[8] (see Figures A1 and A2). This did not provide the same flexibility or range of movements possible by a real human (multi-jointed) neck, but allows the robot to express simple head gestures, such as shaking (side to side), nodding (up and down) and tilting (head to one side).

The head also provided another three DoFs for the eyes: eyes up/down, left/right, and eyelids open/close (Figure A1, A3). Miniature video cameras were also mounted in the eyes (Figure A4). Another two DoFs actuated the mouth; mouth open/close, and mouth smile/sad.

The video cameras incorporated into Kapar's eyes are miniature-type cameras, both with a 1/4 inch B & W CMOS Image sensor producing a PAL output of 288 (H) $\times$ 352 (V) with an effective resolution of 240 TV lines, 1/50 to 1/6000 shutter speed, sensitivity of 0.5 lux/f 1.4. The physical dimensions are approximately 20 $\times$ 14 mm (excluding lugs) with a depth of 25 mm and a weight of approximately 25 g. Three wire connections are available: red = +ve (DC 9 to 12 V, 20 mA max), black = common Gnd and yellow = video out.

The head frame was constructed mainly from sheet aluminium, with custom-machined components produced for the universal joint at the neck. The individual parts are bolted together with machine screws and nuts. The RC servos used were mounted on the head frames by means of screws, and transmission of actuation to the neck, face and eyes achieved by means of push-rods (Figure A5).

All the wiring to the servos used the standard three-wire RC connectors and extensions. The video camera wiring was made using fine twin (+Vs and signal) core screened (0 V) flexible cables. Strain relief for the wires was made at the neck joint by means of cable ties (Figures A2 and A5).

Arms design and construction

The arms were constructed from standard kit parts, which are now available to hobbyists at a reasonable cost for making directly driven joint and link chains from standard size RC servos. The forearms from the original shop dummy were mounted on 6 mm machine screws, and attached to form the hand end of the arms (Figure A6). The shoulder ends of the arms were mounted on plates bolted into the shoulders of the shop dummy (Figure A7). The arm wiring consisted of standard RC three-wire connections from each servo back to the controller board, with strain relief provided by cable ties at appropriate points.

Controller

The controller interface board used is a LynxMotion SSC32 Servo controller board (cf. LynxMotion (2007), Figure A8, right), with the ability to control up to 32 servos simultaneously. Only 16 servos are used for KASPAR's movements, so there is the possibility to use additional servos in future enhancements. The board interfaces the host computer via an RS232 serial port, which is mostly not provided as standard on most modern PCs or laptops. Therefore, a small RS232 to USB adaptor board is also included inside

[8]In fact, the neck joints would normally be described as pan, tilt and yaw. However, because of the unusual configuration of KASPAR's neck linkage, the configuration could be more correctly described as one pan, and left and right compound tilt/yaw movements.

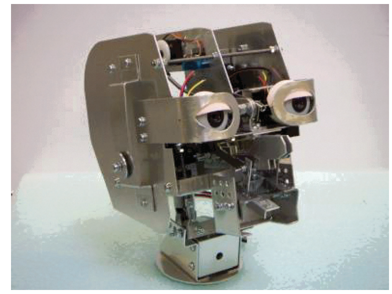

Figure A1. KASPAR's head with silicon rubber face mask removed.

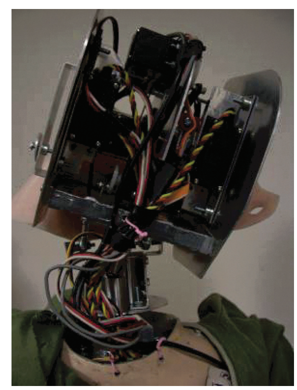

Figure A2. Rear view of head showing wiring and neck joints.

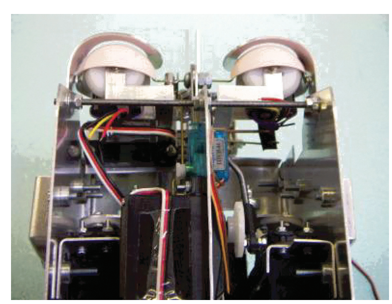

Figure A3. Detailed view of eye actuator linkages.

Figure A4. Miniature video cameras are fitted in each eye.

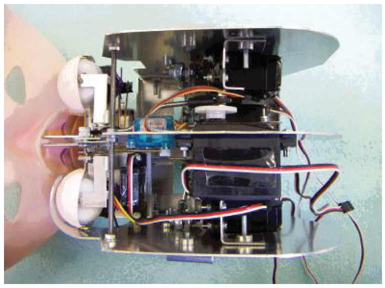

Figure A5. Top view of head showing actuator transmission linkages and wiring.

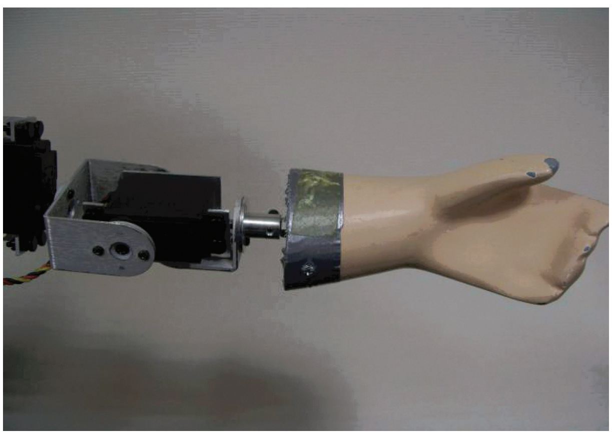

Figure A6. View of arm showing original shop dummy hand attachment.

KASPAR to provide a standard USB interface. Both controller and adaptor board are mounted on an aluminium backplate, which also provides mechanical protection and access to KASPAR's internal systems (Figure A8, left).

Power

For safety, KASPAR is run from two low-voltage lead acid gel batteries. The servo actuators are powered from 6 V, 4 AH battery, and the controller, logic and cameras are supplied by a smaller 12 V, 1 AH battery. Both batteries are protected from short circuits and overload by in-line slow-blow fuses. The main 6 V power fuse

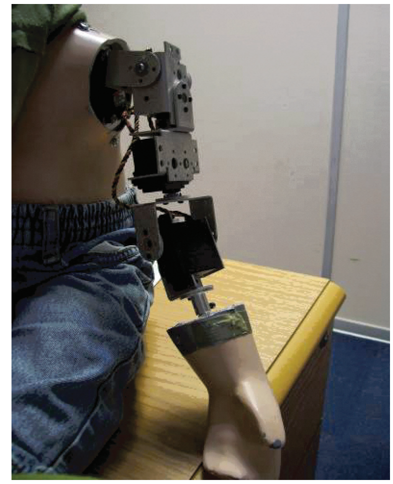

Figure A7. Arm is attached at shoulder end by plates fixed in the dummy body.

value has been set deliberately low (15A) to avoid overloading and subsequent burnout of the expensive high torque shoulder RC servos when manhandling by clients occurs. The batteries are re-charged by two separate chargers, which are connected to KASPAR's batteries by different styles of plug to ensure correct connection. The 6 V charger does not have the capacity to keep the main motor power battery fully topped up while the robot is being used intensively, but if left connected while in use does increase the working time of the robot from about one hour to one and a half hours.

Appendix B

KASPAR II

KASPAR II uses colour video cameras, otherwise the specification of the cameras is identical to those used for KASPAR I, except they are slightly larger with dimensions of 25 × 15 mm and a depth of 20 mm.

KASPAR II's arms use five (one extra over KASPAR I) RC servos apiece, as each incorporates an extra wrist (twist) DoF. The arm links and fittings are custom-made from 1.5 mm thick aluminium sheet, which produces a cleaner, standardised design, avoids the sharp edges which are a feature of the kit linkage parts for KASPAR I and also incorporates extra brackets to mount additional joint position sensors.

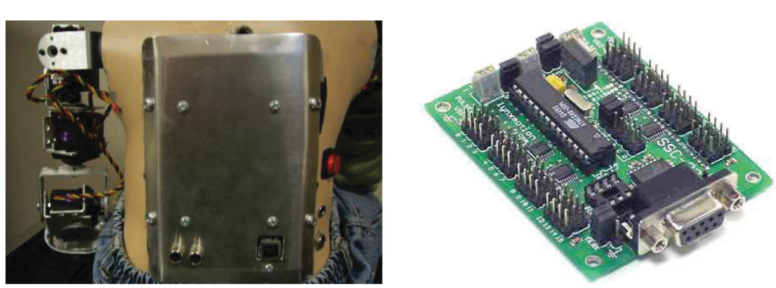

Figure A8. Rear view of KASPAR showing the back plate cover (left); controller board (right).

KASPAR II has arm joint position sensors that provide real-time feedback of the arm positions to the control computer/programme. This is achieved by mounting standard 10 K Ω rotary potentiometers on each arm joint, providing a 0–5 V DC analogue signal proportional to the respective arm joint positions. The analogue signals are then converted by two 8-channel Phidgit USB analogue to digital converters (ADC) which incorporate USB 'pass through' connectors allowing them to be 'daisy-chained' directly onto the standard USB bus to the host control computer.

Since the original dummy body used for KASPAR I was temporarily no longer available for purchase, the dummy body used for KASPAR II is one modelled after a larger approximately six-year-old child, which provides more physical space to accommodate the extra sensors and interface electronics. There is a hole in the chest, with suitable brackets where a Swiss Ranger 3000 (SR3000) general purpose range imaging camera may be mounted enabling straightforward measurement of real-time depth maps. It uses the on-board power supply (12 VDC, 1A max) and interfaces to a host computer via a mini USB 2.00 connection. The specifications are as follows: 176×144 pixels, field of view $47.5^\circ \times 39.6^\circ$, range up to 7.5 m (for 20 MHz modulation), lens f/1.4, illumination power (optical) = 1 W (average power) at 850 nm and physical dimensions $50 \times 67 \times 42.3$ mm (aluminium).

The head mechanism is identical to that used for KASPAR I, although the wiring is made through connectors to allow easy removal and servicing.

Upgrades, changes and planned future improvements

The limited time for which KASPAR can be operational between re-charges has been a problem, and it is desirable to increase the main 6 V battery life. This could be achieved in a number of ways, either by reducing the power requirements of the robot, or by using a larger capacity 6 V battery or charger. Currently, the fuses are mounted internally and require removal of the back plate for access. A relatively simple change would be to use panel-mounted fuses with external access, which would allow operators to change the fuse easily. A more long-term solution would be the incorporation of a flexible current limiting circuit.

While speech interaction is not the main focus of our research, other future applications would like to incorporate speech synthesis, which could be achieved using a dedicated speech synthesis module via the on-board USB adaptor. An on-board microphone for recording interaction partners' speech and sound would be convenient, but the noise generated by the robot may make usage difficult. Both these functions may also be achieved very simply by the incorporation of a loudspeaker and a microphone on the robot which could then be connected to the host computer soundcard input and output connections.

The seven different types of servos originally used have now been standardised to just three types. The four shoulder servos and the base neck servo are high torque types (HiTec 645 MG) and typically cost three times as much as the same sized servos (HiTec HS-422) used for other joints. A single small micro-sized servo (a Supertec NARO HPBB) is used for the pan movement of the eyes. The main limitations with regard to using these RC servos are the relatively poor accuracy obtained and the lack of control feedback. These deficiencies have now been remedied to some extent in new generation servos aimed specifically at the hobby robotics market, but these were not available when KASPAR was designed and built. A review of these new servo types would probably allow the replacement of the original servos with more capable ones, though it is likely that they would be more expensive and require some redesign of the head and arm parts.

The SSC-32 controller has the capacity to potentially control another 16 RC servo actuators. This might be used to add additional facial expressions, or leg movements (for gestures only rather than locomotion in order to maintain the simplicity of the design).

Learning behavior for a social interaction game with a childlike humanoid robot

Frank Broz, Hatice Kose-Bagci, Chrystopher L. Nehaniv, and Kerstin Dautenhahn
Adaptive Systems Research Group
School of Computer Science
University of Hertfordshire, UK
Email: f.broz, h.kose-bagci,c.l.nehaniv,k.dautenhahn @herts.ac.uk

I. Introduction

Social interaction plays a vital role in a child's development. At an early age, children acquire basic social skills, such as mutual gaze and turn-taking, that serve as a scaffold for more sophisticated forms of socially-mediated learning, including language. Children learn these skills through interaction with their caretakers, motivated by intrinsic social drives that cause them to seek out prolonged social engagement as a fundamentally rewarding experience. It is desirable to have robots learn to interact in a similar manner, both to gain insight into how social skills may develop and to achieve the goal of natural human-robot interaction. This paper describes a system for the learning of behavior sequences based on rewards arising from social cues, allowing the iCub, a childlike humanoid robot with a developmentally-inspired design, to engage a human participant in a social interaction game.

II. Social Cues

A. Visual Attention

Gaze is a powerful social cue. It is also one that becomes socially significant at an early developmental stage; even young infants are responsive to other's gaze direction [1]. The simplest gaze cue, and one that is the basis and developmental precursor to more complex gaze behaviors such as joint attention, is the recognition having another's visual attention [2]. This ability is crucial in a social context, as it provides valuable feedback about whether one is interacting with (or has the potential to interact with) someone or whether you are disengaged and merely sharing the same space. While there has been work done in the field of robotics using human gaze patterns to reproduce natural-appearing gaze in robots [3], [4], there has been no work explicitly modeling the role of gaze as a form of social feedback that may guide the robot's overall behavior. In this system, a gaze tracker worn by the human participant is used to collect gaze direction data in real time as sensor input for the robot (and a potential source of reward).

B. Turn-taking

Turn-taking plays a fundamental role in regulating human-human social interaction and communication whereby role-switching and the dynamics are not determined by external forces but emerge from the interaction. It has vital implications in many areas like robot-assisted therapy, especially in studies related to children with autism, where turn-taking games have been used to engage the children in social interaction [5]. Turn-taking is a skill that children begin to develop early in life. Caretakers teach infants how to engage in turn-taking through interacting with them [6]. The cues that regulate turn-taking are multimodal, and may be either general or task-based.

We hypothesize that fluid turn-taking requires attention to the recent history of both one's own and the other's actions in order to anticipate and prepare for the shift in roles. In light of this, the robot's control architecture incorporates a short term memory over the recent history of sensor data relevant to the regulation of turn-taking (to be described in Section IV). Two forms of non-verbal turn-taking are supported in this interaction, drumming and peek-a-boo. Drumming allows the human and robot to engage in turn-taking with clearly defined and easily detectable beginnings and endings. Studies on emergent turn-taking in a drumming interaction have been carried out by the authors previously using a similar childlike humanoid robot [7]. Peek-a-boo is more ambiguous (given sensing limitations) but well understood by human participants, and has also been studied before in embodied human-robot interaction [8].

III. The Interaction History Architecture

This research extends past work on the iCub using the Interaction History Architecture (IHA). IHA is a system for learning behavior sequences for interaction based on grounded sensorimotor histories. While the robot acts, it builds up a memory of past "experiences" (distributions of sensors, encoders, and internal variables based on a short-term temporal window). Each experience is associated with the action the robot was executing when it was recorded, as well as a reward value based on properties of the experience. These experiences are organized for the purpose of recall using information distance as a metric. As the robot acts, the most similar past experience to its current state is found, and new actions are probabilistically selected based on their reward value. For a full description of the architecture, see the journal article [8].

IV. Short Term Memory

In addition to a dynamic memory of sensorimotor experience and associated rewards, it is also useful to have a more detailed, fully sequential memory of very recent experience. This is especially true for skills such as turn-taking, where

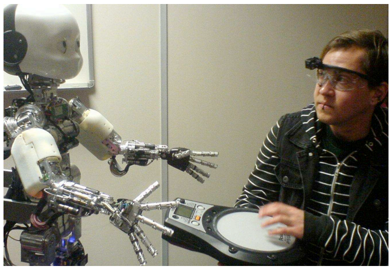

Fig. 1. A person interacts with the iCub using the gaze tracker and drum.

the recent history of relationships between one's own and another's actions must be attended and responded to. While the experience metric space preserves some ordering of experiences (so that rewards over future horizons may be computed), there is not a mechanism to recall the most recent experience, only the most similar. Additionally, experiences aggregate data over a window of time, eliminating potentially useful fine-grained information about changes in sensor values. The short term working memory preserves temporal information about the sensorimotor data over the span of several past experiences. This is especially important for guiding social interactions as it allows rewards to be designed based on these histories of interaction, rather than just the instantaneous state of the interaction that the robot is currently experiencing.

V. Developing Social Interaction

In order to demonstrate these concepts, rewards based on social drives are designed to influence the development of behaviour in an open-ended face-to-face interaction game between the iCub and a human. This work is an extension of an earlier application of IHA to the learning of the game peek-a-boo on the iCub to allow more types of interaction and social cues. The human participant interacts with the robot and may provide it with positive social feedback using their presence and gaze direction, as well as by playing a drum. The rewards representing these social drives for human presence, visual attention, and synchronized turn-taking may be based on the current state of the robot's sensorimotor experience or on the history of experience represented in its short term memory. The robot uses this feedback to acquire behavior that leads to sustained interaction with the human.

The robot comes to associate sequences of simple actions and gestures, such as waving or hitting the drum, executed under certain conditions with successful interaction based on its past experience. The pre-defined actions that the robot chooses among may be either low-level motions that don't have meaning except as part of a sequence or movements that have an (implicit) goal, such as an arm motion to hit the drum. There is no distinction in how these actions are represented internally to the architecture, and both kinds of actions may make up parts of learned behavior sequences. The extended Interaction History Architecture is intended to support the robot developing different socially communicative, scaffolded behaviours in the course of temporally extended social interactions with humans by making use of social drives and its own first-person experience of sensorimotor flow during social interaction dynamics.

VI. Future Work

The role of learning in this system is currently restricted to learning the experience space and the associations between experiences and rewards in order to find effective behavior sequences . But there are many opportunities to extend the role of learning in this system to produce social behavior in a more developmentally plausible manner. While the relationships between sensors monitored for feedback about turn-taking were predefined in this case, one could instead use statistical methods to discover which sensor channels are associated and predictive of one another. This would allow for task-specific turn-taking cues to be discovered, as well as general task-invariant cues. And while gaze is used in this system as a form of social feedback, the robot has no active gaze behavior. It would be interesting to learn action sequences for gaze behavior as well, especially gaze used to regulate turn-taking.

There is also the opportunity to engage in meta-learning about the learned behavior sequences. The topology of the experience metric space could be used to make generalizations about experiences, using the clustering of experiences to identify closely related behaviors or interactions. The ability to make aggregate representations of these clusters that capture their fundamental properties could reduce the computational cost of finding similar experiences, while possibly allowing for more powerful predictions based on current experience, opening up the potential of anticipating the actions of others or even recognizing intent.

Acknowledgment

This research was conducted within the EU Integrated Project RobotCub (Robotic Open-architecture Technology for Cognition, Understanding, and Behaviours) and was funded by the European Commission through the E5 Unit (Cognition) of FP6-IST under Contract FP6- 004370.

References

[1] S. M. Hains and D. W. Muir, "Infant sensitivity to adult eye direction." *Child development*, vol. 67, no. 5, pp. 1940–1951, October 1996.

[2] A. Frischen, A. P. Bayliss, and S. P. Tipper, "Gaze cueing of attention: visual attention, social cognition, and individual differences." *Psychological bulletin*, vol. 133, no. 4, pp. 694–724, July 2007.

[3] Y. Yoshikawa, K. Shinozawa, H. Ishiguro, N. Hagita, and T. Miyamoto, "The effects of responsive eye movement and blinking behavior in a communication robot," in *Intelligent Robots and Systems, 2006 IEEE/RSJ International Conference on*, 2006, pp. 4564–4569.

[4] B. Mutlu, J. Forlizzi, and J. Hodgins, "A storytelling robot: Modeling and evaluation of human-like gaze behavior," in *Humanoid Robots, 2006 6th IEEE-RAS International Conference on*, 2006, pp. 518–523.

[5] B. Robins, K. Dautenhahn, R. te Boekhorst, A. Billard, S. Keates, J. Clarkson, P. Langdon, and P. Robinson, "Effects of repeated exposure of a humanoid robot on children with autism," in *Designing a More Inclusive World*. Springer-Verlag, 2004, pp. 225–236.

[6] D. R. Rutter and K. Durkin, "Turn-taking in mother-infant interaction: An examination of vocalizations and gaze." *Developmental Psychology*, vol. 23, no. 1, pp. 54–61, 1987.

[7] H. Kose-Bagci, D. S. Syrdal, K. Dautenhahn, and C. L. Nehaniv, "Drum-mate: Interaction dynamics and gestures in human-humanoid drumming experiments," *Connection Science*, 2010 (In press).

[8] N. A. Mirza, C. L. Nehaniv, K. Dautenhahn, and R. te Boekhorst, "Grounded sensorimotor interaction histories in an information theoretic metric space for robot ontogeny," *Adaptive Behavior*, vol. 15, no. 2, pp. 167–187, 2007.