RobotCub.org

ROBotic Open-architecture Technology for Cognition, Understanding and Behavior

Cogsys
Cognitive Systems

Project No. 004370

RobotCub
Development of a Cognitive Humanoid Cub

Instrument: Integrated Project
Thematic Priority: IST - Cognitive Systems

D6.2 Results from Interaction Studies on Gesture Communication
WP6 - Gesture Communication

Due date: **01/09/2007**
Submission Date: **26/09/2007**

Start date of project: **01/09/2004** Duration: **60 months**

Organisation name of lead contractor for this deliverable: **University of Hertfordshire**

Responsible Person: **Kerstin Dautenhahn**

Revision: **1.1**

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)	

Contents

1 Executive Summary

WP6 continues to focus on interaction dynamics of social interaction during robot-human play and the prerequisites for gesture and non-verbal communication between robots and humans, as well as the realization of these capabilities in a robot. User studies are used to gain understanding of the kinesics and dynamics of social interaction during robot-human play and its development in ontogeny. At the same time, techniques for achieving this capability in an autonomous robot through grounded sensorimotor experience and interaction histories, are investigated.
This deliverable brings together three areas of complementary research work. The first, defined by two papers detailing an investigation of a robot-human drumming interaction game: "Drum-mate: A Human-Humanoid Drumming Experience" (Kose-Bagci et al., submitted 2007b) and "Emergent Turn-Taking in Drumming Games with a Humanoid Robot" (Kose-Bagci et al., submitted 2007a), both have been submitted to conferences. The results from the drumming interaction showed differences in both the participants qualitative experience and objective measures of the interaction under different conditions of concomitant non-verbal gestures by the robot as well as revealing interesting gender differences.
The second area of work regards an architecture by which a robot can ontogenetically develop through social interaction and grounded sensorimotor experience. We present work that has been published in a journal article (June issue of Adaptive Behaviour (Mirza et al., 2007)) detailing the architecture and experiments using the early interaction game, "peekaboo", between a robot and human. The interaction history was shown to be capable of supporting development of a turn-taking interaction in a robot which took appropriate actions or gestures based on its own grounded sensorimotor experience. Furthermore, we describe in the body of this deliverable key areas of research with the interaction history architecture including issues of scalability, forgetting and experimental technique. Finally, we present the current state of research that brings together the interaction history architecture on a humanoid platform to play the early social interaction game, "peekaboo".
The third area is concerned with the neural bases of gesture communction. Studies during the third project year on this topic include work on (a) eye-contact (work in progress), (b) brain areas involved in gestural communication, and (c) a behavioral study of cooperation and competition during human interaction.

2 Timing and Nonverbal Cues in Interactive Play with a Humanoid Robot

2.1 Gestures in Human-Humanoid Drumming

The results of a study focusing on interaction dynamics of social interaction during human-robot play are presented. The study is an exploratory investigation of a drumming experience between Kaspar, a humanoid robot, and human partners. The social interaction was mediated through a drumming call-and response game and was systematically modulated by non-verbal gestures and cues. The results were statistically analysed in terms of the game performance as well as the evaluation of the game by the participants.
The analysis found a clear effect due to the concomitant gestures during the interaction and also found interesting gender differences between the participants in terms of how they interacted with the robot under different gesture conditions. Male participants tended to focus more on the interaction rather than the accuracy of drumming as gestures increased.

Full details of the experiment and results are to be found in the paper included as *Appendix A*.

2.2 Kinesics of Interaction and Emergent Turn-Taking

In a further series of experiments, we studied emergent turn-taking while regulating the manner in which the robot's actions were produced. In this work, KASPAR uses different probabilistic models to decide when to start and stop its turn. KASPAR uses the number of beats of human participants and the time duration of its previous play as parameters in the models to decide the start and stop times. Therefore the number of beats and play times are not deterministic but emerge completely from the interaction between the human participant and the humanoid. Thus, during the games, sometimes KASPAR plays the 'leader' role, sometimes it follows the human participant. Analysis of the results showed an impact of the turn-taking model on the structure of the interaction in terms of duration and complexity of drumming by human participants as well as on their enjoyment of the interaction game; however, individual differences between participants played a strong role. Moreover participants behaviour changed over the course of (order controlled) exposure to the models, indicating that they may have adapted their interaction to perceived capabilities of the robot.
Full details of the experiment and results are to be found in *Appendix B*.
Both of these results together sugest that deeper study of human-robot interaction kinesics and recipient design is warranted in the area of ontogenetic robotics where a robot develops by engaging in and sustaining social interaction with human partners.

3 Grounded Sensorimotor Interaction Histories and Metric Space of Experience

Over the last three years of the project work has been ongoing to develop the capability for a robot to develop in ontogeny through interaction in a social environment. The history of interaction is composed of grounded sensorimotor experiences which are related to one-another in a metric space defined by their distances in terms of an information theoretic quantity termed Experience Distance.
As a robot interacts with its environment it accumulates experiences and the metric space grows accordingly. A further dimension to the experience is added in terms of rewards signals from the environment. An action selection mechanism can then use this space of experience to decide the next action according to proximity of the current experience to one in its history and the expected value of the experience in terms of reward.
The combination of the sensorimotor experiences, the metric space of experience and the action selection mechanism are collectively referred to as the Interaction History Architecture.

3.1 Peekaboo and the Temporal Horizon of Interaction

Appendix B presents a published article Mirza et al. (2007) that details the architecture as outlined above, and also presents results of robotic experiments that establish the predictive efficacy of the metric space and shows the robot developing the capacity to play the simple interaction game "peekaboo". A quantitative investigation of the appropriate horizon length of experience for the game reveals the relationship between length of experience and cycle time of interaction, and suggests the importance of multiple, and possibly self-adaptive, horizon lengths.

3.2 Technical Assessment of the Interaction History Architecture: Scalability

Appendix C and *D* present recent research on the interaction history architecture. *Appendix C* examines the issue of scalability of the architecture, addressing issues such as forgetting and computational complexity. *Appendix D* utilizes the architecture in a simple simulated toy-problem, in order to study and improve the learning and developmental capabilities of the architecture before implementation on a humanoid robot.

3.3 Current research

Experiments using the peekaboo task with a humanoid robot (KASPAR) will be the focus of the next phase of research on the interaction history architecture and its role in a robot's ontogenetic development of capability to act in a social interaction. The experiments will be based on the previous peekaboo work but will utilize a more generic motivation system, use audio as a further interaction modality as well as utilizing the enhancements learnt through the scalability and simulated test-bed studies. An important goal in this work in particular is to move towards software that will form part of the iCub open software delivered at M42.

4 Neuroscience of Gesture Communication

Work on the neural bases of gesture communication in year 3 was focused on studies in three areas: (a) eye-contact (work in progress), (b) the brain areas involved in gestural communication (Fadiga et al., 2006), and (c) the behavioral study of cooperation and competition during human interaction.

4.1 Eye Contact

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. In addition to this study (preliminary data were presented in, Fadiga, L. Craighero, L., Lungu, O, and Rustichini, A. Eye-to-eye communication, 2005, Society for Neuroscience Meeting, Washington DC), we more recently carried out a behavioral experiment aiming at investigating the gaze behavior of two human subjects while they look each other into the eyes. The parameters we acquired where: eye position (60 Hz), pupil diameter (as an index of attention) and blinking. The basic experimental condition was contrasted with two control conditions. In the first, subjects were looking at themselves through a mirror, in the second, they were looking at a photograph of two eyes. We are currently analyzing the data and we will soon publish these results together with those of the fMRI experiment described above.

4.2 Brain Areas Involved in Gestural Communication

The recent finding that Broca's area, the motor center for speech, is activated during action observation lends support to the idea that human language may have evolved from neural substrates already involved in gesture recognition. Although fascinating, this hypothesis has sparse demonstration because while observing actions of others we may evoke some internal, verbal description of the observed scene. Here we present fMRI evidence that the involvement of Broca's area during action observation is genuine. Observation of meaningful hand shadows resembling moving animals induces a bilateral activation of frontal language areas. This activation survives the subtraction of activation by semantically equivalent stimuli, as well as by meaningless hand movements. Our results demonstrate that Broca's area plays a role in interpreting actions of others. It might act as a motor-assembly system, which links and interprets motor sequences for both speech and hand gestures. These results have recently been published in Social Neuroscience Journal (Fadiga et al., 2006) (*Appendix F*).

4.3 Cooperation and Competition in Human 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 (FDI) 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 subsamples (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 subsample 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.

References

L. Fadiga, L. Craighero, M. Fabbri-Destro, L. Finos, N. Cotillon-Williams, A. Smith, and U. Castiello. Language in shadow. *Social Neuroscience*, 1:77–89, 2006.

Hatice Kose-Bagci, Kerstin Dautenhahn, and Chrystopher L. Nehaniv. Emergent turn-taking in drumming games with a humanoid robot. In *submission*, submitted 2007a.

Hatice Kose-Bagci, Kerstin Dautenhahn, Dag Sverre Syrdal, and Chrystopher L. Nehaniv. Drum-mate: A human-humanoid drumming experience. In *submission*, submitted 2007b.

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*, 15(2): 167–187, 2007.

Appendices

A Drum-mate: A Human-Humanoid Drumming Experience

B Emergent Turn-Taking in Drumming Games with a Humanoid Robot

C Grounded Sensorimotor Interaction Histories in an Information Theoretic Metric Space for Robot Ontogeny

D Technical Report on Interaction History Architecture and Scalability

E Technical Report on the Assessment of the Interaction History Architecture in a Simple Test-bed

F Social Neuroscience Journal Article: Language in Shadow

Drum-mate: A Human-Humanoid Drumming Experience

Hatice Kose-Bagci, Kerstin Dautenhahn, Dag Sverre Syrdal, and Chrystopher L. Nehaniv

Abstract—**We present an exploratory study investigating a drumming experience with Kaspar, a humanoid child-sized robot, and a human. In this work, our aim is not to have Kaspar just replicate the human partner's drumming, but to engage with the human in a 'social manner' using head gestures in a call and response turn-taking interaction and to assess the impact of non-verbal gestures on the interaction. Results from the first implementation of a human-robot interaction experiment are presented and analysed qualitatively (in terms of participants' subjective experiences) and quantitatively (concerning the drumming performance of the human-robot pair). The interaction experience is discussed in terms of imitation, turn-taking, and the effect of gender differences.**

***Index Terms*—Humanoid, robot drumming, human-robot interaction, imitation**

I. Introduction

MUSIC performance is a good tool for studying the interaction between humans and robots in terms of social aspects including imitation, turn-taking and synchronization. Drumming is one of the best ways of performing music in robotics, since it is relatively straightforward to implement and test, and can be implemented technically without special actuators like fingers or special skills or abilities specific to drumming.

There are several works concerning music performance in human-robot interaction. In [1,2,3] robotic percussionists, play drums in collaboration with human partners. These artifacts are robot arms connected to upper torsos that are specially designed to play drums. In [4], an approach based on the movement generation using dynamical systems was tested on a Hoap-2 humanoid robot using drumming as a test case. Similarly, in [5] humanoid drumming is used as a test bed for exploring synchronization.

However, a robot will also need to motivate and sustain drumming behaviour coping with a wide range of users. One way of motivating such behaviour is through the use of social gestures. 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 long-term interactions [6,7], in maintaining user involvement with the tasks encouraged by the agent.

Applied to the field of robotics, the need for the possession of a set of social skills for a robot in order to encourage behaviour successfully may require that it possesses the ability to use social cues and gestures to motivate users to interact with it. This is especially the case for assistive robotics [8].

We can already find robotic systems that use social gestures in order to encourage human-robot interaction. 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 [9]. Other recent examples include small cartoon like robotic "creatures" such as Keepon and Roillo designed to be used in interaction with children [10,11]. These little rubber 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 as in the Wizard of Oz (WoZ) technique, as a part of social interaction. Other related work is discussed in section 2.

In this study, our humanoid robot Kaspar 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 which are meant to motivate the human. Kaspar produces fixed head gestures and eye-blinking as it drums. Our approach is tested with adult participants in several drumming sessions, and the experimental results are reported and analyzed below in terms of imitation, turn-taking, and the impact of gender differences.

The rest of this paper is organized as follows; in the next section, related research is summarized. Next, the methodology is briefly described. The section 4 presents the research questions, corresponding achievements and conditions. The experiments are described in the section 5. Section 6 includes a brief conclusion on what was learned from this work, and the final section presents ideas for future work.

II. Related Research

HAILE [1,2] is a robot arm that aims to play a drum in collaboration with a human partner to study social, mathematical, physical, and technological aspects of music. HAILE does not use fixed deterministic rules, but uses

H. Kose-Bagci, K. Dautenhahn, D. S. Syrdal, and C. L. Nehaniv are with the Adaptive Systems Research Group, University of Hertfordshire, School of Computer Science, Hatfield, Herts, AL 10 9AB, U.K.
(e-mail: {h.kose-bagci,k.dautenhahn,d.s.syrdal, c.l.nehaniv}@herts.ac.uk).

autonomous methods to create variant rhythms. It perceives a variety of complex features of the human partner's drumming whereby a microphone on the drum analyses the sounds and produces rhythms in response.

In [3], a somewhat 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 discover the right tempo, and it trains itself. It uses a simple threshold mechanism to understand 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.

ROILLO is a simple robot with a rubber coated foam head, body, and an antenna. It has 3 wires connected to simple servos which move the head and body in various directions. It is used in experiments with children to study interactions between robot and children [10].

KEEPON is another simple robot, which has only a rubber head and a body. It is has a small CCD camera, and microphone on it. It can move its head, turn its body, and make bobbling 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 [11,12].

III. Methodology

In the current study the human partner plays a rhythm which Kaspar tries to replicate, in a simple form of imitation (mirroring). 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 (deterministic) 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 with the minimum values allowed by its joints. Kaspar needs at least 0.3 seconds 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 Fig. 1, the basic model of Kaspar-human interaction is presented. The model requires the gestures of both human and the humanoid for social interaction, as well as drumming. Currently human gestures are not detected and therefore excluded from the current implementation.

One of the fundamental problems in this scenario is the timing of the interaction; timing plays a fundamental role in the regulation of interaction (cf. [13]). It is not always clear when the robot or human partner should start interaction in taking a turn. Currently, in this model some predefined fixed time duration heuristics are used for synchronization. Kaspar starts playing if the human partner is silent for a few seconds, and tries to motivate the human partner with simple gestures.

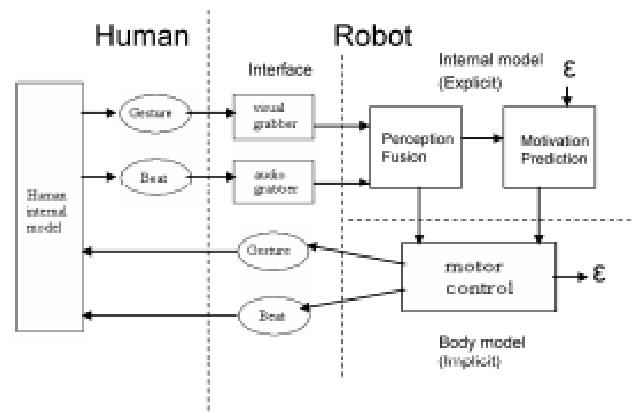

Fig. 1 The model for Kaspar-human interaction

IV. Research Questions, Corresponding Achievements and Conditions

In this work, the effect of the robot’s social gestures in a game of imitation, and turn-taking, was studied. A simple drumming game enriched with the robot’s gestures was used as a test bed, and the subjective evaluations of the participants were analysed. Our primary focus was the possible impact that utilizing social gestures would have, not only on the game itself (in terms of performance), but also on the participant’s subsequent evaluation of the game.

Fig. 2 Snapshots of Kaspar’s dynamic head gestures used in the experimental tests

We studied three conditions with increasing amounts of gesturing. In the first condition Kaspar does not use any gestures. Kaspar only imitates the drumming. This condition was called *no-gesture*. In the second condition, simple head gestures (e.g. moving the head to the right or left, moving the head up or down, tilting head slightly to different angles) and eye blinking were included in Kaspar's movements (Fig. 2). Kaspar starts drumming with one of the fixed gestures. If the human partners do not play their turn, then Kaspar does not do anything, too, and the turn passes to the partner. A fixed order of *n* gestures was used, and this order was repeated for every *n* turns. The value for *n* should be large enough so that the human partner does not realize that this is a fixed pattern but rather that the gestures are either meaningful or random (In the current experiments n was 7). This condition is called *gesture* in our experiments. In the third condition, Kaspar simply repeats the sequence of gestures without playing even if the partners did not play their turn. This case is named *gesture+*. The gestures and their sequences were the same in the last two conditions, and the drumming part was the same in all of the three conditions.

V. EXPERIMENTS

A. Kaspar

The experiments were carried out with the humanoid robot called Kaspar. Kaspar is a child-like humanoid robot which was designed and built by the members of the Adaptive System Research Group at the University of Hertfordshire to study human-robot interactions with a minimal set of expressive robot features. Kaspar has 8 degrees of freedom in the head and neck and 6 in the arms and hands. The face is a silicon-rubber mask, which is supported on an aluminum frame. It has 2 DOF eyes fitted with video cameras, and a mouth capable of opening and smiling, see description in [14].

B. Experimental Setup

The experiments were carried out in a separate room isolated from other people and noises which could affect the drumming experiment. Kaspar was seated on a table with the drum on its lap. The human partner was seated in front of the robot using another drum that was fixed on the table (Fig. 3). The human participants used a pencil to hit the drum. Although we suggested to 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.

Fig. 3 A screen shot from the experiments

C. Software Features

The implementation of robot perception and motor control used the YARP environment [15]. YARP is an open-source framework used in the project RobotCub that supports distributed computation that emphasizes robot control and efficiency. It enables the development of software for robots, without considering a specific hardware or software environment. Portaudio [16] software was used to grab audio from the audio device, within the YARP framework.

The acoustic sound waves recorded by the sound grabber module are converted to digital music samples, which allows to use 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.

D. Participants

Six female participants in the age range of 21-66, and six male participants in the age range of 24-30 took part in the study. All participants were right-handed and worked in computer science or similar disciplines at the University. They had not interacted with Kaspar prior to the experiment, and they were overall not familiar with robots. None of our participants had children, expect for one participant who had grown up children and grandchildren.

E. Interaction Game Setup

We used a one minute demo of the robot without any game where participants were shown how to interact with Kaspar. This was followed by three games reflecting the three experimental conditions described above each lasting three minutes, without indicating to the participants anything about the differences between the conditions. We used all six possible different presentation orders of the games, to analyze the effect of the order of the games on the humans. To account for possible fatigue or learning by the participants, in the *sequential order* section, we analyse 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, e.g. for one participant the first game (number 1) would be the no-gesture game, and for another participant, it would be the third game (number 3).

F. Evaluation of Questionnaire Data

After the experiment the participants were asked to complete a questionnaire investigating their preferences and opinions on the three experimental conditions.

1) **Most and least preferred game types:**

The frequencies of participants which rated each game as most preferred can be seen below in Table 1.

TABLE 1
MOST PREFERRED GAME

Game type	Participants
no-gesture	2
gesture	6
gesture+	3
No preference	1

Table 1 shows that the most popular game type was the *gesture* game, while *no-gesture* and *gesture+* type were less preferred.

The frequencies of participants which rated each game as least preferred can be seen in Table 2.

Table 2 shows that no participants considered the *gesture* game as the least preferred, while the *no-gesture* and *gesture+*

game, had a similar number of participants which considered them the least preferred.

TABLE 2
LEAST PREFERRED GAME

Game type	Participants
no-gesture	6
gesture	0
gesture+	5
No preference	1

2) Gender Differences in most and least preferred game types:

Most and least preferred game type according to gender are described below in Table 3 and Table 4 and in figures 4 and 5.

TABLE 3
MOST PREFERRED GAME ACCORDING TO GENDER

			Gender		
			Male	Female	Total
Game Type Most Preferred	nogesture	Count	2	0	2
		Expected Count	.9	1.1	2.0
		Std. Residual	1.1	-1.0	
	gesture	Count	2	4	6
		Expected Count	2.7	3.3	6.0
		Std. Residual	-.4	.4	
	gesture+	Count	1	2	3
		Expected Count	1.4	1.6	3.0
		Std. Residual	-.3	.3	
Total		Count	5	6	11
		Expected Count	5.0	6.0	11.0
		Std. Residual			

TABLE 4
LEAST PREFERRED GAME ACCORDING TO GENDER

			Gender		
			Male	Female	Total
Game Type Least Preferred	nogesture	Count	1	5	6
		Expected Count	2.7	3.3	6.0
		Std. Residual	-1.0	1.0	
	gesture+	Count	4	1	5
		Expected Count	2.3	2.7	5.0
		Std. Residual	1.1	-1.0	
Total		Count	5	6	11
		Expected Count	5.0	6.0	11.0
		Std. Residual			

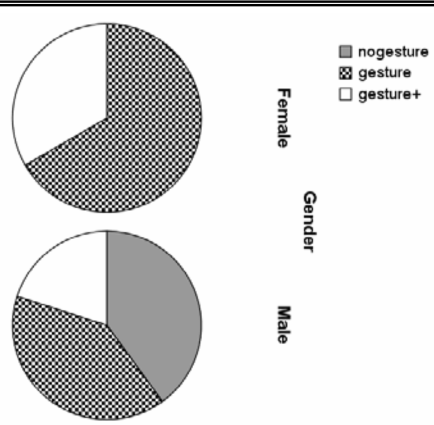

Fig. 4 Most preferred game type according to gender

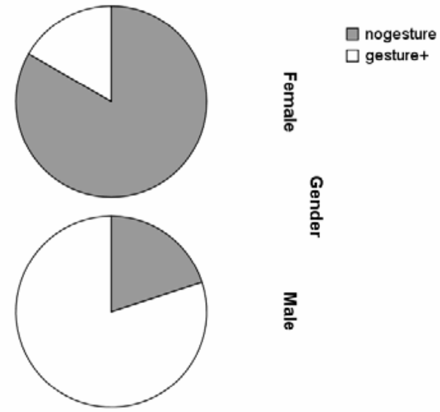

Fig. 5 Least preferred game type according to gender

The differences between males and females in most preferred game type seem to indicate that for males the preferred game type is spread evenly across the three game types, while for females the *no-gesture* game is not preferred by any. The differences between males and females were assessed using a likelihood ratio chi-square test. This test found no significant differences due to gender for this sample size (($\chi^2(1,11)=3.70, p= .16$).

For the least preferred game type, however, there were significant differences due to gender (($\chi^2(1,11)=4.75, p= .03$). As table 4 and figure 5 suggests, this difference manifests as males predominantly choosing the *gesture+* game type as their least preferred game type, while females predominantly chose the *no-gesture* game type as their least preferred game.

3) Reasoning behind preferences

While an exhaustive description of the qualitative analysis of the participants' responses is beyond the scope of this brief paper, a short description will be given below:

Two main themes emerged from the analysis, reflecting two different ways of evaluating the games. The first theme was that of *task-based evaluation*, in which participants would explain their choice by referring to the success of Kaspar in imitating their drumming. The second theme was that of *interaction-based evaluation*, wherein participants would explain their choice as to which games they preferred the most and least by referring to their enjoyment of the interaction and their general liking for the robot.

When the results from the qualitative analysis were compared to the preferences of the participants, it was clear that the task-based evaluation led participants to rate the *gesture+* evaluation as their least favorite game type. Participants using an interaction-based evaluation would choose the *no-gesture* game type as their least favorite.

In terms of gender differences, more female participants used an interaction-based evaluation when explaining their preferences compared to male participants. There were, however, some males who used the interaction-based evaluation. As such, differences between the males and females in this sample may reflect a greater tendency in males to use a task-based evaluation when evaluating the game types.

G. *Behavioural Data*

1) **Sequential order**

The error is the average difference between the human's number of beats and Kaspar's number of beats in each turn. It was observed that the average error in the number of beats decreased inversely to the sequence number in order of the games presented. The participants usually 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. Interestingly, without any external encouragement, as they got used to the game, they progressively were able to synchronize themselves to the robot better. The number of errors decreased significantly between the first and third trials ($Z=2.275, p<.05$). Details of the results are presented in Table 5.

TABLE 5
OBSERVED BEHAVIOUR ACCORDING TO ORDER

Order	Avg. error	Max # of beats	Avg. # of beats	Avg. # of turns
1	4.1 ± 3.6	41	7.5 ± 5.3	17.2 ± 6.2
2	3.1 ± 3.4	37	6.2 ± 4.3	18.3 ± 6.7
3	2.3 ± 1.8	16	4.6 ± 2.2	20.5 ± 3.8

2) **Interaction game type**

The *gesture* game had the highest average error, followed by the *gesture+* game. The non-gesture game had the smallest error rate. However, the differences between games were not significant in this sample size ($Z=1.01, p=.27$).

TABLE 6
OBSERVED BEHAVIOUR ACCORDING TO GAME TYPE

Game type	Avg. error	Max # of beats	Avg. # of beats	Avg. # of turns
no-gesture	2.6 ± 2.5	41	6.2 ± 4.6	18.6 ± 5.8
gesture	3.5 ± 3.8	37	6.2 ± 5.0	18.8 ± 5.7
gesture+	3.3 ± 3	31	6.0 ± 4.7	18.9 ± 6.4

The maximum number of beats decreased with the increasing amount of gestures in the game. The average number of beats also slightly decreased with the increasing amount of gestures in the game. The average number of turns was almost the same in all three games. The total number of beats tend to decrease as the amount of gestures in the game increased. Again, the differences between game type were not significant in this sample size. Details are presented in Table 6.

3) **Gender**

While the sample size makes it difficult to make any strong inferences as to differences between groups, it may be reasonable on the basis of our analysis to present the results from each gender separately. Our qualitative analysis suggests that there are differences in the dynamics when interacting with the robot and as such presenting the results from the male and female sample separately, rather than just focusing on the differences, may be more informative to the reader. However, the small number of participants makes inferential statistics problematic, and as such the following analysis is only descriptive.

Sequential order

In terms of order, in later games, the participants tended to have more turns with fewer beats, which helped them synchronize with Kaspar better - decreasing Kaspar's error rate in drumming and increasing the success of the interaction (Table 7 and Table 8).

TABLE 7
OBSERVED BEHAVIOUR OF MALES ACCORDING TO ORDER

Order	Avg. error	Max # of beats	Avg. # of beats	Avg. # of turns
1	4.1 ± 3.3	41	7.8 ± 6.0	15.0 ± 6.5
2	3.6 ± 3.9	37	7.3 ± 4.5	15.0 ± 5.0
3	2 ± 0.7	11	4.5 ± 0.9	19.0 ± 1.6

TABLE 8
OBSERVED BEHAVIOUR OF FEMALES ACCORDING TO ORDER

Order	Avg. error	Max # of beats	Avg. # of beats	Avg. # of turns
1	4.2 ± 4.3	36	7.2 ± 5.0	19.0 ± 5.7
2	2.6 ± 3.2	31	5.1 ± 4.2	22.0 ± 6.4
3	2.5 ± 2.5	16	4.8 ± 3.1	21.7 ± 5.1

Game type

For male participants, the total number of beats decreased with the increasing amount of gestures in the games. These results suggest that, as the number of Kaspar's gestures increased, they tended to focus more interaction and less on than drumming (Table 9).

TABLE 9
OBSERVED BEHAVIOUR OF MALES ACCORDING TO GAME TYPE

Game type	Avg. error	Max # of beats	Avg. # of beats	Avg. # of turns
no-gesture	2.5 ± 2.7	41	7.3 ± 5.8	16.2 ± 6.1
gesture	3.9 ± 3.6	37	6.3 ± 4.8	16.3 ± 5.8
gesture+	3.0 ± 2.9	9	5.7 ± 2.6	15.8 ± 4.2

Different from the male participants, for females the total number of beats increased as the amount of gestures in the game increased. This indicates that the female participants tended to become more involved in the drumming with increases in non-verbal interaction gestural cues. The detailed evaluations are presented in Table 10.

TABLE 10
OBSERVED BEHAVIOUR OF FEMALES ACCORDING TO GAME TYPE

Game type	Avg. error	Max # of beats	Avg. # of beats	Avg. # of turns
no-gesture	2.7 ± 2.5	16	5.0 ± 3.0	21.0 ± 4.8
gesture	3.1 ± 3.8	36	5.8 ± 5.0	21.3 ± 4.7
gesture+	3.5 ± 3	31	6.3 ± 4.7	20.3 ± 7.8

It is interesting to note that, although the error rate in *gesture+* was less than in the *gesture* condition, male participants liked it the least overall. They thought too many gestures distract them from drumming, instead of enjoying the gestures. Although *gesture* had the worst error rate, overall they liked it the most. In contrast, although the error rate in *gesture+* was the highest, female participants liked it more than the *no-gesture* game which had the lowest error rate (Table 10).

VI. CONCLUSIONS

In this work, we introduced a computational model of an imitative rhythmic interaction game using non-verbal gestural head, neck, and blinking gestures and deterministic turn-taking between a robot and human partner. We based our model on drumming, which is a very suitable task for testing human-robot interaction. It is intended as more than a simple drumming synchronization task. In the long-term, we aim to develop social interaction between the robot and the human partner, which would not simply focus on synchronization to produce the same tempo, but result in producing a joyous and fruitful experience, while allowing us to gain insight into the role of non-verbal gesture in sustaining and regulating human-robot interaction.

We used drumming interaction games enriched with different amounts of Kaspar's gestures to motivate the humans. In our experiments, we saw that humans are, in fact motivated by gestures and take enjoyment from this sharing. Too many gestures, however, break their concentration. Drumming with no gestures is considered successful by participants in terms of a drumming task but it is not considered successful in terms of social interaction. The results from this experiment thus highlight the possible tradeoff between the participants' subjective evaluation of the drumming experience, compared to objective measurements of the drumming performance, also reflecting individual preferences as to task and interactional aspects of the task. These results point towards a clear role for the use of appropriate amount and types of non-verbal gestures as a means of motivating drumming behaviour and regulating the interaction when interacting with a robot.

The reason for the high error rates at the start of the games is probably due to the human partner's high expectations from the game. Especially the male partners appeared to view this experiment not as a game, but rather a task to complete. Also, due to their background the human partners might have tried to 'test' the robot's limitations. So they initially played very fast, and very long sequences, and used different parts of the drum to enrich their play. They expected Kaspar to watch, understand and imitate them (most of the human partners thought the robot could detect them with its eye cameras and that the gestures were meaningful). As they played more, they understood the limited capabilities of the robot and modified their drumming and tried to synchronize with it.

Both the female and male participants overall liked the games with gestures, which had the worst error rates in the evaluations. This shows that the right amount of gestures would attract their attention, and make their experience enjoyable, although it did not actually help their drumming. This reveals a strong difference between the subjective evaluations of the drumming experience by the participants, compared to objective performance measures.

This work is a first step in human-robot interaction research on synchronization, timing, and turn-taking using drumming games. Although we started with a simple implementation, the results are unexpected and interesting. As explained above, in our setup Kaspar just repeated the beats produced by the human partner, and made simple fixed head gestures accompanying its drumming (we especially used very simple gestures, not complex ones like smiling or frowning in order not to affect the human participants too much). The human partners' in return, perceived these simple behaviours as more complex and meaningful. They adapt themselves to the system unconsciously.

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, and adapted themselves to the game better, so the success rate improved over time. Humans, as shown here, were not passive subjects in this game, but adapted themselves unconsciously to the capabilities of the robot. In order to facilitate and motivate such adaptation, aspects of the interaction that are not directly related to the task itself, such as interactional gestures may play an important role.

VII. FUTURE WORK

Based on these results, future work on the humanoid drumming system will involve further study of the use of gestures for motivating the human partners. Because of our promising results from using gestures, we foresee a system wherein Kaspar 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 show satisfaction with these interactions. If our initial results can be extrapolated, then such a system will be even more capable of motivating and sustaining 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 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.

REFERENCES

[1] G. Weinberg, S. Driscoll and 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, pp. 456-461.

[2] G. Weinberg, and Driscoll, S. (2006). Robot-human interaction with an anthropomorphic percussionist. *Proceedings of International ACM Computer Human Interaction Conference (CHI 2006)*. Montreal, Canada, pp.1229-1232.

[3] C. Crick, M. Munz, and Scassellati, B. (2006). Synchronization in social tasks: Robotic drumming. *Proceedings of IEEE RO-MAN 2006*, pp. 97-102.

[4] S. Degallier, C. P. Santos, L. Righetti and A. Ijspeert. (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)*.

[5] S. Kotosaka, and Schaal, S. (2000). Synchronized robot drumming by neural oscillator. *International Symposium on Adaptive Motion of Animals and Machines.*

[6] T. W. Bickmore & Cassell, J. (2005). Social Dialogue with Embodied Conversational Agents. In: J. v. Kuppevelt, L. Dybkjaer & N. Bernsen (Eds.), *Natural, Intelligent and Effective Interaction with Multimodal Dialogue Systems*. New York: Kluwer Academic, pp. 23-54.

[7] T. W. Bickmore & Picard, R. W. 2005. Establishing and Maintaining Long-Term Human Computer Relationships. *ACM Transactions on Computer-Human Interaction* 12(2): 293-327.

[8] A. Tapus & Mataric', M. J. (2006). Towards Socially Assistive Robotics. *International Journal of the Robotics Society of Japan* 24(5), pp.576-578

[9] C. Breazeal (2002). Designing Sociable Robots. MIT Press.

[10] M.P. Michalowski, S. Sabanovic and Michel, P. (2006). Roillo: Creating a social robot for playrooms. *Proceedings of IEEE RO-MAN 2006*, pp.587-592.

[11] H. Kozima, C. Nakagawa, Y. Yasuda, and Kosugi, D. (2004). A toy-like robot in the playroom for children with developmental disorder. *Proceedings of the International IEEE Conference on Development and Learning (ICDL-04, San Diego, USA)*.

[12] KEEPON. (2007) http://univ.nict.go.jp/people/xkozima /infanoid/robot-eng.html#keepon.

[13] B. Robins, K. Dautenhahn, C. L. Nehaniv, N. A. Mirza, D. Francois, and Olsson, L. (2005). Sustaining interaction dynamics and engagement in dyadic child-robot interaction kinesics: Lessons learnt from an exploratory study. In *Proc. of the 14th IEEE International Workshop on Robot and Human Interactive Communication, ROMAN2005*, pp. 716-722.

[14] M.P. Blow, K. Dautenhahn, A. Appleby, C. Nehaniv, D. Lee (2006). Perception of robot smiles and dimensions for human-robot interaction design. *Proc. 15th IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN06),* pp.469-474.

[15] G. Metta, P. Fitzpatrick, and Natale, L. (2006). Yarp: yet another robot platform. International Journal on Advanced Robotics Systems, Special Issue on Software Development and Integration in Robotics, 3(1), pp. 43-48.

[16] Portaudio (2007). http://www.portaudio.com/trac/wiki/.

[17] H. Kose, and Akin, H. L. (2001). Object recognition in robot football using a one-dimensional image. The Tenth Turkish Symposium on Artificial Intelligence and Neural Networks (TAINN 2001), pp.291-300.

Emergent Turn-Taking in Drumming Games with a Humanoid Robot

Hatice Kose-Bagci
Adaptive Systems Research Group
School of Computer Science
University of Hertfordshire
+44 1707281031
h.kose-bagci@herts.ac.uk

Kerstin Dautenhahn
Adaptive Systems Research Group
School of Computer Science
University of Hertfordshire
+44 1707284333
k.dautenhahn@herts.ac.uk

Chrystopher L. Nehaniv
Adaptive Systems Research Group
School of Computer Science
University of Hertfordshire
+44 1707284470
c.l.nehaniv@herts.ac.uk

ABSTRACT

We present an exploratory study investigating emergent turn-taking in a drumming experience involving Kaspar, a humanoid child-sized robot, and adult participants. In this work, 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. Therefore the robot is not just 'following' and imitating the human, but could be the leader in the game and being imitated by the human. Results from the first implementation of a human-robot interaction experiment are presented and analysed qualitatively (in terms of participants' subjective experiences) and quantitatively (concerning the drumming performance of the human-robot pair).

Keywords

Humanoid, emergent turn-taking, robot drumming, human-robot interaction, imitation.

1. INTRODUCTION

Turn-taking is an important ingredient of human-human interaction and communication whereby role switch ('leader' and 'follower') is not determined by external sources but emerges from the interaction. Human beings '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 novel framework which enables emergent turn-taking, and role-switching between a human and a humanoid in an imitation game.

There are several example works that studied turn-taking in games and conversations in the literature, focusing on different aspects. In [1] there is a very good example of emergent turn-taking from the domain of developmental psychology. This work gives the example of the emergent turn-taking between a mother and a baby without any explicit control mechanism. 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 sucking actions.

One of the most difficult issues in teaching and education of children with autism is to teach children the concept of 'turn-taking'. In [2,3] turn-taking games are used to engage children with autism in social interactions. Another example of turn-taking games is given from a cognitive robotics view in [4]. 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 alteration. But here the turn-taking behaviour was led by the human experimenter in reaction to the robot's visually driven actions.

Ito and Tani studied joint attention and turn-taking in an imitation game played with the humanoid robot QRIO, where the human participants try to find the action patterns, which were learned by QRIO previously, by moving synchronously with the robot [5].

From a linguistics point of view, in [6], some of the important features of turn-taking in human conversation identified are:

1. Speaker-change recurs, or at least occurs.
2. Mostly, one party talks at a time.
3. Occurrences of more than one party speaking at the same time are common but brief.
4. Transitions (from one turn to the next) with no gap and no overlap are common (slight gap or slight overlap is accepted).
5. Turn order is not fixed, but varies.
6. Turn size is not fixed, but varies.
7. Length of conversation is not specified in advance.
8. What parties say is not specified in advance.
9. Relative distribution of turns if not specified in advance.
10. Number of parties can vary.
11. Talk can be continuous or discontinuous.

Built on these features, Thorisson developed a turn-taking mechanism for conversations based on his previous work on the *Ymir mind model* for communicative creatures and humanoids [7]. The expressive humanoid robot KISMET [8,9] used social cues for regulating turn-taking in non-verbal interactions with people. Here, a sophisticated robot control architecture modeling motivations, emotions and drives was used 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. Similarly, in our work we study emergent turn-taking, but based on minimal, probabilistic control models.

Our particular test bed for studying emergent turn-taking here is human-robot drumming games. We used imitation games involving drumming as a test bed since they seem a suitable tool for studying the interaction between humans and robots in terms of social aspects including imitation, turn-taking and synchronization. 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. Drumming is relatively straightforward to implement and test, and can be implemented technically without special actuators like fingers or special skills or abilities specific to drumming [10]. There are several works concerning drumming in human-robot interaction. In [11,12,13] robotic percussionists play drums in collaboration with human partners. These artifacts use robotic arms that are specially designed to play drums. In [14], an approach based on the movement generation using dynamical systems was tested on a Hoap-2 humanoid robot using drumming as a test case. Similarly, in [15] humanoid drumming is used as a test bed for exploring synchronization.

In this study, our humanoid robot Kaspar plays 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. With a simple, but novel probabilistic method 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 our previous work [10] where we tested deterministic turn-taking. In this work Kaspar used no gestures, but only drumming to interact with the human. We found in our previous work that different robot nonverbal gestures influence people's responses in the drumming game, and thus decided to carry out this experiment without any gestures in order to be able to focus our analysis on the turn-taking behaviour.

The rest of this paper is organized as follows: the next section describes the methodology. Section 3 presents the research questions and expectations. The experiments, results and analysis are described in the section 4. Section 5 includes a conclusion on what was learned from this work, and presents ideas for future work.

2. METHODOLOGY

In the previous study [10] the human partner played a rhythm which Kaspar tried to replicate, in a simple form of imitation (mirroring). Kaspar had 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, due to its limited motor skills. It tailored the beats beyond its skills with the minimum values allowed by its joints: Kaspar needed at least 0.3 seconds between each beat to get its joints 'ready', so that even if the human plays faster, Kaspar's imitations would be slower using durations of at least 0.3 seconds between beats. It also needed to wait for a few seconds before playing any rhythm in order to get its joints into correct reference positions.

One of the fundamental problems in this scenario is the timing of the interaction; timing plays a fundamental role in the regulation of interaction (cf. [16]). It is not always clear when the robot or human partner should start interaction in taking a turn. Therefore, in the previous work, some predefined fixed time duration heuristics were used for synchronization. Kaspar started playing if the human partner was silent for a few seconds, and tried to motivate the human partner with simple nonverbal gestures.

In this work, we used a probability based novel approach for timing and turn-taking. These emerge from the interaction between the human and the humanoid. Three simple models are used to control the start and stopping of the robot's regular drumming based on the duration time of the previous turn and number of beats played in the previous turn for the interaction partners. In this work, we will simply name the models as *model1*, *model2* and *model3*. The *model1* is a step function, *model2* is a simple triangular function and *model3* is a parabolic function (see Appendix A). The outputs are limited by maximum and minimum limits to ensure that Kaspar and the human would have time to play at least once in every turn. There is a minimum time threshold of 1.5 seconds (experimentally determined) for human participants. Notice that this minimum time threshold is the input to the model, there is no threshold to the output of the model. So the minimum time per turn could be smaller than 1.5 seconds. For Kaspar the minimum time threshold is one second and the beat threshold is one beat. These minimum values are assigned to the output of the model if it is below these values, so the minimum values could not be smaller than these values. The only model which does not have threshold limitations is *model3* due to its parabolic nature (Appendix A (c),(f)). In every turn, Kaspar looks up the probability of start or stop, and takes action accordingly. For the start Kaspar uses the time duration of its last play, and for the stop, the number of beats of the human participant from the previous turn is used.

The human will start the game and Kaspar will use its turn-taking strategy, when the human participant stays silent for two seconds (only for the first turn). After the first turn, the turn-taking strategy is always determined by Kaspar's models. The computational models are presented in Equations (1), (2), and (3).

$$p(x) = \begin{cases} 0, x < Th \\ 1, x >= Th \end{cases} \qquad (model1) \qquad (1)$$

$$p(x) = x / Th \qquad (model2) \qquad (2)$$

$$p(x) = 1 - 1/x \qquad (model3) \qquad (3)$$

where *Th* represents the threshold of time in *start* (Appendix A (a),(b)) and beat (Appendix A (d),(e)) in *stop*. (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.) As a function of the previous duration and number of beats in the interaction, according to their respective probability functions (1), (2), (3), the three models may return value 1, which triggers starting or stopping in the turn-taking algorithm:

```
Algorithm turn-taking
   ThTime_i= KasparTime_i-1
   IF model_j (HumanTime_i,ThTime_i) = 1
      THEN KASPAR STARTS
   ThBeat_i= HumanBeat_i
   IF model_j (KasparBeat_i,ThBeat_i) = 1
      THEN KASPAR STOPS
```

So at every turn, Kaspar decides when to start and stop according to the performances of both the human participant, and itself. Thus, the game is not deterministic but emerges from the current status from both Kaspar and the human participant.

3. Research Questions & Expectations

In this work, the effect of the different computational models on emergent turn-taking in an imitation game was studied. A simple drumming game enriched with different models determining the turn-taking strategy of the humanoid robot was used as a test bed, and the subjective evaluations of the participants were analysed. Our primary research questions were:

1) How do different robot turn-taking strategies based on computational, probabilistic models impact the drumming performance of the human-robot pair?

2) How do the different robot turn-taking strategies impact the participants' subjective evaluation of the drumming experience?

We expect to have 'successful' games in terms of turn-taking emerging from the social interaction between human and the humanoid. Our 'success' criteria would be the number of turns with no or slight overlaps and gaps. Also the number of human beats detected by the robot and number of beats played by the robot itself would give us hints about the quality of the games.

We did not include any head or body gestures other than drumming to observe the impact of the models clearly. We also set up simulated experiments before the real experiments, to define the maximum and minimum limits and thresholds for the real experiments with humanoid and human participants.

We studied three models with different parameters. Each model is used both for starting and stopping the robot's play. 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 threshold. As described in the previous section in detail, *model1* was a step function, where the new value of could not be 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 parts cannot play longer than 10 seconds per turn). Unlike *model1*, *model2* has a triangular shape which has the threshold as an upper bound. 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 human participants, so we foresee that the model would not be as popular as the other two models among the participants. The last condition is *model3* which is a parabolic model which cannot be bounded by the thresholds. It reaches high values (close to 1) very fast compared to *model2*. Therefore we predict that it would give more play time and enable to play more beats than *model2*. Also, in our simulations we noticed that it could enable good games (with a very low number of overlaps and conflicts) if we played short sequences, but since it is not bounded with thresholds, it 'reacts' to the human but does not exactly 'imitate' the games, which might not be accepted by participants.

4. Experiments, Results & Analysis

4.1 Kaspar

The experiments were carried out with the humanoid robot called Kaspar. Kaspar is a child-like humanoid robot 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 8 degrees of freedom in the head and neck and 6 in the arms and hands. The face is a silicon-rubber mask, which is supported on an aluminum frame. It has 2 DOF eyes fitted with video cameras, eyelids capable of blinking, and a mouth capable of opening and smiling, see description in [17].

Figure 1 A screen shot from the experiments showing a person playing a drumming game with Kaspar.

4.2 Experimental Setup

The experiments were carried out in a separate room isolated from other people and noises which could affect the drumming experiment. Kaspar was seated on a table with the drum on its lap. The human partner was seated in front of the robot using another drum that was fixed on the table (Figure 1). The human participants used a pencil, or their bare hands to hit the drum. Although we suggested to 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.

4.3 Software Features

The implementation of robot perception and motor control used the YARP environment [18]. YARP is an open-source framework used in the project RobotCub that supports distributed computation that emphasizes robot control and efficiency. It enables the development of software for robots, without considering a specific hardware or software environment. Portaudio [19] software was used to grab audio from the audio device, within the YARP framework.

The acoustic sound waves recorded by the sound grabber module are converted to digital music samples, which allows using 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.

4.4 Participants

12 participants in the age range of 23-32 (4 female and 8 male) took part in the study. All participants were right-handed and worked in computer science or similar disciplines at the University. Only two of them had interacted with Kaspar prior to the experiment, and they were overall not familiar with robots. Three of our participants had children aged 1-3 years.

4.5 Interaction Game Setup

We used a one minute demo of the robot without any drumming game involved where participants were shown how to interact with Kaspar. This was followed by three games reflecting the three experimental conditions described above each lasting three minutes, without indicating to the participants anything about the differences between the conditions. We used all six possible different presentation orders of the games, to analyze 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 analyse 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, e.g. for one participant the first game (order 1) would be the *model1* game, and for another participant, *model1* would be the third game (order 3).

4.6 Evaluation of Questionnaire Data

After the experiment the participants were asked to complete a questionnaire investigating their preferences and opinions on the three experimental conditions.

4.6.1 Most and least preferred game types:

The number of participants which rated each game as most preferred can be seen below in Table 1.

Table 1. Most preferred game

Game type	Participants
model1	6
model2	0
model3	6

Table 1 shows that both the *model1* and *model3* games were preferred by the same amount of participants, while no participant most preferred *model2*.

The number of participants which rated each game as least preferred can be seen in Table 2.

Table 2 shows that most of the participants considered the *model2* game as the least preferred, while the *model1* and *model2* games had a small number of participants which considered them the least preferred. The *model3* game was slightly more popular than the *model1* game.

Table 2. Least preferred game

Game type	Participants
model1	3
model2	8
model3	1

4.6.2 Most and least preferred games according to sequential order

The number of participants which rated each game as most preferred according to the sequential order can be seen below in Table 3. It is shown that the most popular game type was the third game, while first and second games were less preferred.

Table 3. Most preferred game

Order	Participants
1	3
2	2
3	7

The number of participants which rated each game as least preferred can be seen in Table 4. All ordinal positions of occurrence in the sequence of the games had a similar number of participants which considered them the least preferred.

Table 4. Least preferred game

Order	Participants
1	4
2	3
3	4

4.6.3 Reasoning behind preferences

While an exhaustive description of the qualitative analysis of the participants' responses concerning their impressions and preferences about the drumming games is beyond the scope of this brief paper, a short description will be given below:

The order of the games had an impact on the participants. Their liking of the games increased significantly between the first and third trials (for drumming, $F(2,22)=3.29$, $p=0.069$; for sociality , $F(2,22)=4.904$, $p<0.05$, with ANOVA). They preferred the last game more, which could be because they got used to the scenario as they played more, so they had more successful plays as they spent more time; this is consistent with our previous findings [10].

According to the game types there appeared also to be an impact (for drumming, F(2,22)=2.444, p=0.110; for sociality, F(2,22)=2.895, p=0.77, with ANOVA). Notice that number of participants is small which makes harder to perform a good statistical analysis.

The participants liked the *model2* game less, because due to the model's nature, it gives the least play time to the human and Kaspar. So Kaspar does not seem to imitate the human participants' game at all, but rather 'plays on its own' (Kaspar plays at least one beat even when it does not detect a response from the human participant). So conflicts occurred between Kaspar's and the human's turns, and Kaspar seemed to take the turn when the human was still playing. Most of the human participants found this annoying and some of them even called this action "rude". The *model1* gives the human participant the most play time, and since it uses the previous play's play time as a threshold, it ensures that the current play time is at least as long as the previous play time. The *model3* is not limited by the thresholds, but its probabilities are increasing fast, so it does not give small values very often, and it mostly yielded turn durations in the 1.1-1.3 seconds range.

Therefore according to the explanations of the human participants in the questionnaires, they liked the *model1, and model3* because, they felt the robot could imitate them better in these games. But some of the participants also mentioned that the *model1* game, was 'slower' than the other games. In fact, since they were given more play time than the other games, there were time gaps between their turns and Kaspar's turns, so they felt the tempo of the game was slower than the others. These participants preferred *model3* which cannot give as much play time for human participants as *model1,* since it does not change with the threshold, but gives a long enough time to have a coordinated game. They mentioned that the tempo of the game was faster than the *model1* game. In this game both human and Kaspar had 3-4 beats every turn, there were less conflicts than when using *model1,* and less gaps compared to *model2* between two games. But due to the nature of the model, Kaspar played similar patterns which seemed to be independent of the human participants' performance, which annoyed some of the participants. Still one participant found this like "teaching her son to play drum". 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".

4.7 Behavioural Data

4.7.1 Sequential order

There is 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 order of presentation as they got used to the scemario as they played more (Tables 5-8).

Table 5. Observed behaviour of human (beats) according to order

Order	#of turns	#of nonzero turns	Max # of beats	Total # of beats (Kaspar's view)	Total # of beats (real)
1	93±45.08	27.83±14.3	5	44.33±25.8	104.3±27.5
2	91.1±43	29±12	4	47.8±27	114.8±34.5
3	90.4±44.27	32.3±15.4	5	55.1±32.69	122.8±23.8

Table 6. Observed behaviour of human (duration) according to order

Order	Avg. time per turn	Max time per turn	Min time per turn	Total time
1	0.99±0.567	3.11	0.01	70±27.53
2	0.99±0.6	2.06	0.01	69.1±27
3	1±0.57	3.11	0.01	68±24.8

Table 7. Observed behaviour of KASPAR (beats) according to order

Order	Avg. # of beats per turn	Max # of beats	Min # of beats	Total # of beats
1	1.7±0.773	6	1	136±31.86
2	1.74±0.76	6	1	136±29.2
3	1.81±0.71	7	1	139±22.9

Table 8. Observed behaviour of KASPAR (duration) according to order

Order	Avg.time per turn	Max time per turn	Min time per turn	Total time
1	1.08±0.11	3	1	97.8±41.4
2	1.07±0.09	3	1	95.5±39.6
3	1.07±0.09	4	1	94.8±41.2

4.7.2 Interaction game type

According to the game types, in *model2* more turns were played and there is a significant difference between the *model2* and the other two models, but the number of nonzero turns (where the human played at least one beat) is less compared to the other two models (Table 9). Although from our observations human participants appeared to play similarly in all three games, Kaspar could detect a very small number of the human participants' beats in case of *model2* (Kaspar only detected human participants' beats, and recorded them, when it decided that the humans' play a turn according to its computational model. Kaspar discarded the beats played by human participants at other times, namely during Kaspar's own play times). The other two models showed almost similar behaviour, with *model1*'s number of turns and number of nonzero turns slightly larger compared to *model3*. When the total number of beats were compared, however, we realize that the number of beats of the human participant is larger (Table 9) and the number of beats of Kaspar is smaller in *model1* than in games using *model3* (Table 11). Likewise. the time given for human to play is longer (Table 10), and the time given to Kaspar to play (Table 12) is shorter in *model1* than in games using *model3*, both in terms of total time and the avg. time per turn. In terms of maximum and minimum durations per turn, for Kaspar (Table 12) there is no significant change, but in the case of the human player both are significantly longer for *model1* than for the other models (Table 10).

So as we observe from tables 9-12, in the *model2* game Kaspar almost played alone, and did not give much chance to the human player in terms of play time. We put a minimum threshold of 1.5 seconds as input to the models, but we did not put any threshold to the output of the models, so *model2* could give very little play time to the human player according to the probabilities. So in this model Kaspar is mainly a leader and not a follower.

In the *model3*, Kaspar is given a chance to be a follower and leader almost equally. Kaspar and the human had almost equal total durations of play (Table 10, and Table 12). Kaspar had more impact on the play and played longer rhythms (Table 11).

In *model1* the human was given more time than Kaspar (Table 10, and Table 12), but Kaspar played more beats than the human participants (Table 9, and Table 11). Whereas in *model3,* Kaspar and human participant where given almost equal durations and opportunities to play.

Table 9. Observed behaviour of Human (beats) according to game type

Game type	#of turns	#of nonzero turns	max # of beats per turn	Sum of beats (Kaspar's view)	Total # of beats (real)
model1	65.1±4.03	37.3±15	5	72.1± 27.8	113±29.223
model2	151±3.46	21.1±7.8	3	25.6±9.67	116.7±24.69
model3	59±1.5	31±13	5	50±22	112.08±34.9

Table 10. Observed behaviour of Human (duration) according to game type

Game type	Avg.time per turn	Max time per turn	Min time per turn	Total time
model1	1.53±0.02	3.11	1.5	99.3±5.31
model2	0.25±0.01	0.61	0.01	37.4±1.7
model3	1.2±0.01	1.8	1	70±1.8

Table 11. Observed behaviour of KASPAR (beats) according to game type

Game type	Avg. beat per turn	max # of beats per turn	Min # of beats per turn	Total # of beats
model1	1.55±0.3	5	1	99.9±11
model2	1±0.01	3	1	153±4.22
model3	2.7±0.1	7	2	158±3.5

Table 12. Observed behaviour of Kaspar (duration) according to game type

Game type	Avg.time per turn	Max time (sec) per turn	Min time per turn	Total time
model1	1±0.04	3	1	67±3.05
model2	1±0.004	3	1	151±3.204
model3	1.2±0.04	4	1	70±1.71

5. CONCLUSIONS

In this work, we introduced three probabilistic computational models of an imitative rhythmic interaction game that facilitates emergent turn-taking between a robot and human partner. We based our model on drumming, which is a very suitable task for testing human-robot interaction. It is intended as more than a simple drumming synchronization task. We aim to develop social interaction between the robot and the human partner, which would not simply focus on synchronization to produce the same tempo, but result in producing a joyous and fruitful experience, emerging from human-robot social interaction.

We used drumming interaction games enriched with different probabilistic computational models which enables Kaspar to start and stop its turns using its observations on the human participant's play. According to the play time per turn and number of beats played during a turn, Kaspar starts and stops its own turn, and therefore influences the human participant's turn. So each turn is emerging from the current play status of Kaspar and human participants. This is more similar to natural human-human conversation, where human beings start and stop their turns in conversations and also in non-verbal communication according to criteria of their own without an external or internal rigid 'clock'.

There are three game models used in this work. *Model1* has a step function, *model2* is a simple triangular function and *model3* is a parabolic function. In the first two models we use the duration of play and the number of beats of previous turns as thresholds. Also we make sure that there is a minimum value of thresholds to have a continuous game, where Kaspar and the human would have enough time to play each turn, or play at least one beat.

We analysed the games in terms of sequence of order, and according to game type. We used the questionnaires and the observed behaviors of the Kaspar and human participants to compare the games.

In terms of sequence there is an impact on the participants drumming behaviour and evaluation of the games while they played the three games. They tend to beat more, in fewer turns, and, in terms of the questionnaire data, they liked the games more as they played more.

In terms of game types, in *model1*, the total number of beats for Kaspar were higher than the total number of beats for the human participants (100/65), whereas, the total game duration is higher for human participants than for Kaspar (70/100). In *model3*, the total number of beats is lower than for *model1*. Although the total play durations for Kaspar and the humans were almost identical, the total number of beats for Kaspar was almost three times as high as that of human participants. In *model1* and *model3* almost half of the turns were nonzero (i.e. the human played at least one beat). The *model2* game has the largest number of turns which is almost twice as high as the other two, but the number of nonzero turns is small (14%). Also, here the robot beats much more than the human (number of human beats is 17% of Kaspar's beats). Therefore Kaspar was mainly a leader in the *model2* game, and did not give much chance to the human participants to play. So when compared to the questionnaire results, it is logical that humans did not like the *model2* game. Since the average time for human participants in *model3* game is smaller it could be viewed as more ‘natural’. But it was not bounded by thresholds so it was not affected by the performance of the human players. There were many overlaps between Kaspar’s play turns and human participants’ play turns in *model2*. So either Kaspar or the human participants interrupted the other which was found ‘annoying’ by the human participants, and caused the loss of detection of human participants’ beats for Kaspar (as described before, Kaspar did not ‘listen’ when it played itself). Although there were gaps between the humans’ and the robot’s turns in *model1*, and *model3* did not seem to imitate the human participants in every turn, both models were successful in terms of emergent turn-taking. Although we used very simple models, and this work is a first step in this domain, we were able to observe some very good games in terms of coordinated turn-taking, and some of the participants even compared the game to a normal game you play with your children.

It is important to note that while Kaspar's drum playing changed in terms of timing based on simple models, some human participants commented that Kaspar behaved intelligently, e.g. they thought that the robot interrupted them in a structured way, in order “to tell them something”. We aimed not to imitate the human participants’ drumming exactly, but tried to get some emergent effects from the interaction between human and humanoid instead. Although some of the participants found this ”annoying” since Kaspar did “not imitate them well”, surprisingly another group of the participants thought Kaspar played like a small child, and they enjoyed the games.

Also over time, the participants learned the limits of Kaspar and the rules of the game, and adapted themselves to the game better, so they had better games, in terms of turn-taking and synchronization. We could observe long sequences of plays without any overlaps or gaps between the turns, and human participants were really enthusiastic about the games. Humans, as shown here, were not passive subjects in this game, but adapted themselves unconsciously to the capabilities of the robot. 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” [21, p. 71]. According to conversation analysis, this turn-taking is integral to the formation of any interpersonal exchange [21, p. 66]. While in our study the robot’s behaviour was controlled based on simple computational models, the human participants used their recipient design skills in the interaction.

The issue of recipient design will be explored further in our future research. Also, we plan to add robot gestures to our future games (using head movements and facial expressions), since most of the participants commented in the questionnaires that gestures might improve Kaspar’s social interaction skills, and we observed the same result in our previous work [10].

Acknowledgments

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 help in preparing the experiments and the paper: Michael L. Walters, Joe Saunders, N. Assif Mirza, Qiming Shen, Aris Alissandrakis, Nuno Otero and Dag S. Syrdal.

References

[1] H. Hendriks-Jansen. Catching Ourselves in the Act: Situated Activity, Interactive Emergence, Evolution, and Human Thought. MIT Press, Cambridge, Mass., 1996.

[2] K. Dautenhahn and A. Billard. Games children with autism can play with Robota, a humanoid robotic doll. In S. Keates, P. M. Langdon, P.J. Clarkson, and P. Robinson, editors, *Universal Access and Assistive Technology*, pages 179-190. Springer-Verlag, London, 2002.

[3] B. Robins, K. Dautenhahn, R. te Boekhorst, and A. Billard. Effects of repeated expoisure of a humanoid robot on children with autism. In S:. Keates, J. Clarkson, P. Langdon, and P. Robinson, editors, *Designing a More Inclusive World*, pages 2255236. Springer Verlag, London, 2004.

[4] R. A. Brooks. Personal Communication at the 2nd International Conference on Cognitive Technology (CT’97), Aizu, Japan, August 28, 1997

[5] M. Ito and J. Tani: "Joint attention between a humanoid robot and users in imitation game", Proc. 3rd Int. Conf. on Development and Learning (ICDL'04), La Jolla, USA, 2004.

[6] Sacks, H., Schegloff, E. A., & Jefferson, G. (1974). A simplest systematics for the organization of turn-taking for conversation. *Language*, 50, 696-735.

[7] Thórisson, K. R. (2002). Natural Turn-Taking Needs No Manual: Computational Theory and Model, from Perception to Action. In B. Granström, D. House, I. Karlsson (Eds.), Multimodality in Language and Speech Systems, 173-207. Dordrecht, The Netherlands: Kluwer Academic Publishers.

[8] C. Breazeal. *Designing Sociable Robots*, The MIT Press, 2002.

[9] C. Breazeal. "Towards sociable robots," T. Fong (ed), *Robotics and Autonomous Systems*, **42**(3-4), pp. 167-175, 2003.

[10] H. Kose-Bagci, K. Dautenhahn, D. S. Syrdal, and C. L. Nehaniv. “Drum-mate: A Human-Humanoid Drumming Experience”, –submitted.

[11] G. Weinberg, S. Driscoll and 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, pp. 456-461.

[12] G. Weinberg, and Driscoll, S. (2006). Robot-human interaction with an anthropomorphic percussionist. *Proceedings of International ACM Computer Human Interaction Conference (CHI 2006)*. Montreal, Canada, pp.1229-1232.

[13] C. Crick, M. Munz, and Scassellati, B. (2006). Synchronization in social tasks: Robotic drumming. *Proceedings of IEEE RO-MAN 2006*, pp. 97-102.

[14] S. Degallier, C. P. Santos, L. Righetti and A. Ijspeert. (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)*.

[15] S. Kotosaka, and Schaal, S. (2000). Synchronized robot drumming by neural oscillator. *International Symposium on Adaptive Motion of Animals and Machines.*

[16] B. Robins, K. Dautenhahn, C. L. Nehaniv, N. A. Mirza, D. Francois, and Olsson, L. (2005). Sustaining interaction dynamics and engagement in dyadic child-robot interaction kinesics: Lessons learnt from an exploratory study. In *Proc. of the 14th IEEE International Workshop on Robot and Human Interactive Communication, ROMAN2005*, pp. 716-722.

[17] M.P. Blow, K. Dautenhahn, A. Appleby, C. Nehaniv, D. Lee (2006). Perception of robot smiles and dimensions for human-robot interaction design. *Proc. 15th IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN06)*, pp.469-474.

[18] G. Metta, P. Fitzpatrick, and Natale, L. (2006). Yarp: yet another robot platform. International Journal on Advanced Robotics Systems, Special Issue on Software Development and Integration in Robotics, 3(1), pp. 43-48.

[19] Portaudio (2007). http://www.portaudio.com/trac/wiki/.

[20] H. Kose, and Akin, H. L. (2001). Object recognition in robot football using a one-dimensional image. The Tenth Turkish Symposium on Artificial Intelligence and Neural Networks (TAINN 2001), pp.291-300.

[21] D. Boden 1994 ***The Business of Talk: Organizations in Action***, Cambridge: Polity Press/Basil Blackwell

APPENDIX A: COMPUTATIONAL MODELS (see text for explanation)

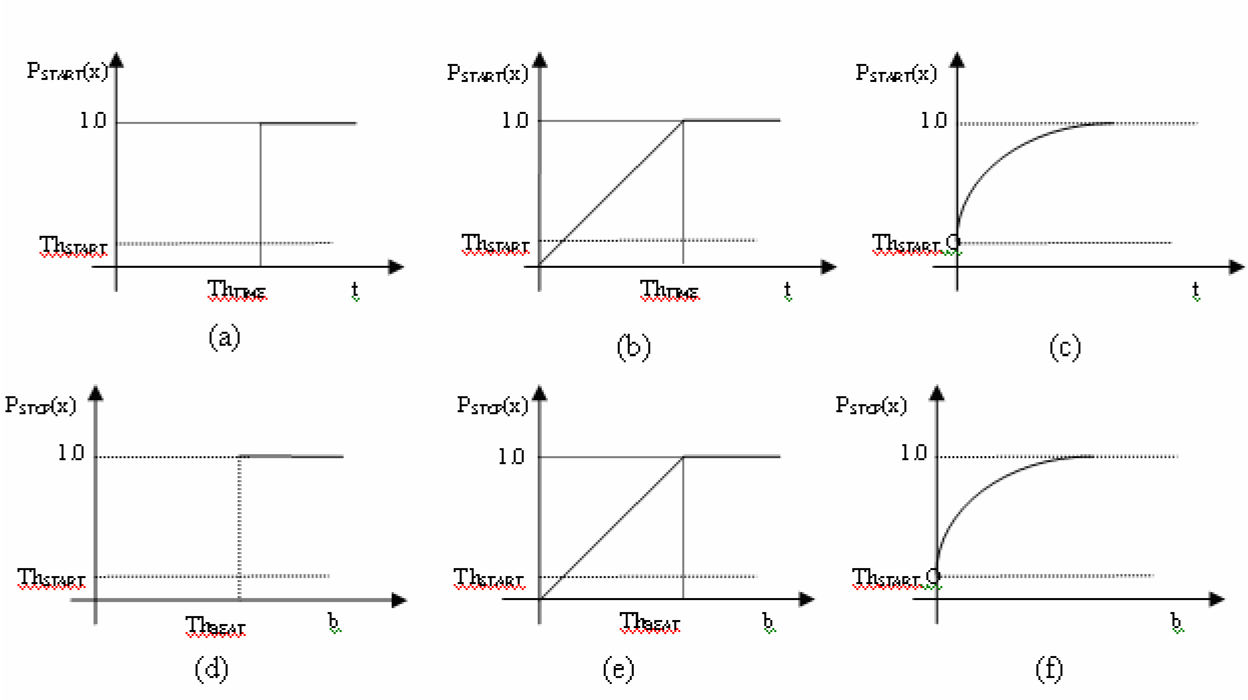

Adaptive **Behavior**

Grounded Sensorimotor Interaction Histories in an Information Theoretic Metric Space for Robot Ontogeny

Naeem Assif Mirza, Chrystopher L. Nehaniv, Kerstin Dautenhahn, René te Boekhorst
Adaptive Systems Research Group, School of Computer Science, University of Hertfordshire, UK

We motivate and present a definition of an embodied, grounded individual sensorimotor interaction history, which captures the time-extended behavior characteristics of humans and many animals. We present an architecture that connects temporally extended individual experience with capacity for action, whereby a robot can develop over ontogeny through interaction. Central to this is an information theoretic metric space of sensorimotor experience, which is dynamically constructed and reconstructed as the robot acts. We present results of robotic experiments that establish the predictive efficacy of the space and we show the robot developing the capacity to play the simple interaction game "peekaboo." A quantitative investigation of the appropriate horizon length of experience for the game reveals the relationship between the length of experience and the cycle time of interaction, and suggests the importance of multiple, and possibly self-adaptive, horizon lengths.

Keywords interaction history · sensorimotor experience · information theory · peekaboo · ontogenetic development

1 Introduction

A challenge for research into embodied cognition in robots is to reach beyond reactive architectures to systems that exhibit the time-extended behavior characteristics of humans and many animals. We are interested in how cognitive structures in natural and artificial systems can arise, which capture the history of interactions and behaviors of an agent actively engaged in its environment, without resorting to ungrounded symbolic representations of past events. Our goal is to design and test such an architecture for a robotic agent, addressing the problem of broadening the temporal horizon to generate adaptive behavior, while not necessarily trying to model details of human behavior. The ultimate aim of the work is to achieve scafolded ontogeny in robots and other artificial agents by endowing them with an extended temporal horizon grounded in their own sensorimotor interaction histories. In this work we lay the theoretical and experimental groundwork for one attempt at achieving this.

We introduce an architecture for ontogeny and adaptive action based on a metric space of temporally extended sensorimotor experience. The robot chooses how to behave in the world based on what it has experienced. This results in further experience modifying the space of experience, establishing a tight coupling of experience and action.

Correspondence to: Naeem Assif Mirza, Adaptive Systems Research Group, School of Computer Science, University of Hertfordshire, College Lane, Hatfield, Hertfordshire, AL10 9AB, UK.
E-mail: N.A.Mirza@herts.ac.uk,
Tel.: +44 1707 28 1133, *Fax*: +44 1707 28 4185

(2007), Vol 15(2): 167–187.
DOI: 10.1177/1059712307078662

In Sections 2 and 3, we establish a theoretical basis for our particular view of an interaction history, including the information theoretical aspects, ending by presenting a computational robotic model. Related research is discussed in Section 4. A simple experiment is presented in Section 5, which demonstrates the efficacy of the space generated by the robot passively experiencing its environment. The architecture is then used by a robot to develop the capacity to engage in peekaboo, a simple early interaction game (Section 6). We conclude with a discussion of the experimental results, the current strengths and limitations of the model, and suggestions for future work.

2 Interaction Histories

We start by considering how memory is viewed from an embodied perspective and why temporal extension is important. We then draw on this motivation to present a suitable definition of interaction history, which can become the basis for our robotic model.

2.1 Temporal Horizon and Extension

The "temporal horizon" of an agent delimits the history (whether personal or socially acquired) that an agent has access to (Nehaniv, Polani, Dautenhahn, te Boekhorst, & Cañamero, 2002). Autonomous embodied artificial agents, which make use of interaction histories in guiding their actions, can be thought of as extending their temporal horizon beyond that of a simple "reactive agent," for instance, Braitenberg vehicles (Braitenberg, 1984). These agents become post-reactive systems when acting with respect to a broad temporal horizon by making use of temporally extended episodes in interaction dynamics (Nehaniv et al., 2002). Internal state, as used in affective agents, can also extend the temporal scope of the agent (potentially indefinitely but usually for the short or medium term), as previous interactions can affect later actions through the agents' affective state. However, in general this approach does not allow for access to episodic historical events and so cannot, for instance, suggest more complex alternative courses of action (Scheult & Logan, 2001).

We note that the temporal horizon for an agent potentially encompasses the entire past history of the agent (although it can be focused on episodes of horizon of arbitrary size). History may inform forward temporal extension in, for example, prediction, anticipation and planning. The size of the temporal horizon influencing behavior can be varied and does vary between natural agents. Some agents, it seems, live only in the present, for instance Braitenberg vehicles[1] and probably bacteria.[2]

Research into the developmental psychology of human infants points to the importance of anticipation and prediction in the development of cognitive capabilities (see, for example, von Hofsten, 1993). A traditional artificial intelligence approach to achieving this might be to build an internal model of the process or task in question, and then to use that model to predict future states. However, we argue that by using a temporally extended history as the basis for action, links between experiences and actions may be built that allow the agent to act such that it exhibits the appearance of prospection, predicting repeated and familiar events in its environment.

2.2 Dynamic Systems, Cognition, and Memory

Cognitive systems can be viewed as the structure and processing of dynamical systems operating in various types of state space (agent–environment, sensorimotor, perception–action, etc.; Dautenhahn & Christaller, 1996; Kelso, 1995; Thelen & Smith, 1994). Regions and attractors (or structures) of these dynamical systems may reflect interesting areas in terms of remembering and adaptive action. These structures are created through interplay of the dynamic system and the agent interaction with the environment.

From an action-oriented point of view, an agent's interaction with the environment can construct the structures that are used for remembering how to act. Furthermore, the process of remembering and acting may alter those structures thus reconstructing the "memory." This may involve altering the detail of the original structures, changing the relative importance of them or, in terms of dynamical systems, moving and altering the attractors. We refer to this process as dynamical construction. To illustrate, consider auto-associative Hopfield artificial neural networks (Gurney, 1997). The dynamics of such networks resolve to particular attractors (memories) on presentation of particular inputs. Learning of new memories affects what is already stored, and if the network were able to learn while recalling, recall would also modify "stored" memories. Thus, memory consists not of static representations of the past that can be recalled with perfect clarity, but rather is the result of a dynamic accretion of interaction with the environment.

2.3 Remembering, Memory, and Action

We follow the argumentation of Rosenfield (1988)—for a review see Clancey (1991)—and Dautenhahn and Christaller (1996) in relation to situated cognition, that human and animal memory is the result of an accumulation of interaction with the environment. Furthermore, the way that memory manifests itself is as embodied action. That is, it is in actions resulting from recall that we witness memory and that recall itself is dependent on embodiment. This argument has support in the view that the purpose of perception and memory for the natural environment is to guide action (Glenberg, 1997) and that even abstract concepts can be interpreted in terms of physical actions and properties.

Glenberg (1997), Clancey (1997), and Pfeifer and Scheier (1999), among others,[3] also argue for an embodied situated memory and memory as recategorization. The emphasis is on the interaction with the environment and a process view of memory.

An important aspect of interaction history is that it is constructed from the perspective of the individual, that is, it is autobiographical in nature. Dautenhahn (1996) defines an autobiographical agent, as "an embodied agent that dynamically reconstructs its individual history (autobiography) during its lifetime" (page 31).

In terms of the accepted separation of memory types due to Tulving (1983), interaction histories could be classified as episodic memory as opposed to semantic memory. That is, it is the memory of events (with a temporal aspect and, usually, a personal aspect), rather than the memory of knowledge and categories. Interaction histories, though, have elements of both. Categories and knowledge may emerge from many overlapping experiences aided by the process of dynamic construction, while certain unique events may still stand out and give memory its episodic nature. This is a view supported by Glenberg (1997).

An autobiographical agent may also be able to communicate significant episodes in its past to other agents, which could further increase the temporal horizon of the agent and that of others (Nehaniv, 1999). Here, the notion of recounting, or communication of, that history is important particularly in social agents.

While we do not claim that an interaction history can describe all aspects of (human) memory, we believe that exploring its features may give insights into the nature of memory in adaptive behavior as a whole.

2.4 Ontogenetic Development

Ontogenetic development in artificial and natural organisms can be seen as an incremental, possibly open-ended, self-organizing process of change, where an organism refines its current capabilities by using internally generated drives and motivations and exploration of its environment and embodiment to generate new goals, capabilities and behaviors (Lungarella, Metta, Pfeifer, & Sandini, 2003).

We hypothesize that a dynamically constructed history of interactions, which is used to generate and select actions in an embodied agent, can serve to scaffold the ontogenetic development of the agent. Development in this case can be seen as the increasing richness of the connections of experience with action, mediated by suitable mechanisms. Such a history can facilitate incremental development at the borders of experience. It is known that this is the case for human development, which is continually scaffolded by building new capabilities on top of existing ones. Learning proceeds at the periphery of known experience and already mastered interaction skills enabling development ("zone of proximal development;" Vygotsky, 1978).

However, the development process depends on drives and motivation. Classical conditioning and two-process reinforcement learning based on positive and negative reinforcers (e.g., Rolls, 1999) are potential mechanisms for connecting previous experience with choice of action. In this study, an internally generated motivation system (see Appendix A) is used that assigns reinforcement values to an episode of experience.

2.5 Definition of an Interaction History

In view of the preceding discussion and motivated by a dynamical systems, embodied view of memory, we propose the following definition of an interaction history as being:

> the temporally extended, dynamically constructed, individual sensorimotor history of an agent situated and acting in its environment including the social environment, that shapes current and future action.

The key aspects of this definition are the following.

- Temporal extension. The overall horizon of an agent's experience extends into the past (potentially including all previous experience available to the agent) and also into the future in terms of prediction, anticipation, and expectation.
- Dynamical construction. This indicates that the history is continually being both constructed and reconstructed, with previous experiences being modified in this process, and potentially affecting how new experiences are assimilated.
- Grounding. The history need not be symbolic (i.e., recorded in terms of externally imposed representations) and is grounded in the sensorimotor experience of the agent. Beyond innate structures for perception, any new representations and categories may emerge in cognitive structures as a result of the agent–environment interaction.
- Remembering in action. The process of remembering drives and shapes the choice of current and future action, while dynamically reshaping the structures employed in remembering.

Note that we use the term "interaction" to indicate that this temporally extended history encompasses the sensorimotor history, the history of action as well as the feedback of action on the history. This definition encompasses all types of interaction with the environment, but specifically includes the social environment. It differs from simple reinforcement or neural net learning in explicitly incorporating the temporally extended nature of experience.

3 An Interaction History Architecture

Figure 1 shows an architecture that demonstrates how histories of sensorimotor experiences can be explicitly integrated into the control of a robot. Our approach is to continually gather sensorimotor data and find episodes of sensorimotor experience in the history near to the current episode and, depending on the course of subsequent experience, choose from among actions that were executed when these episodes were previously encountered, or possibly other actions.

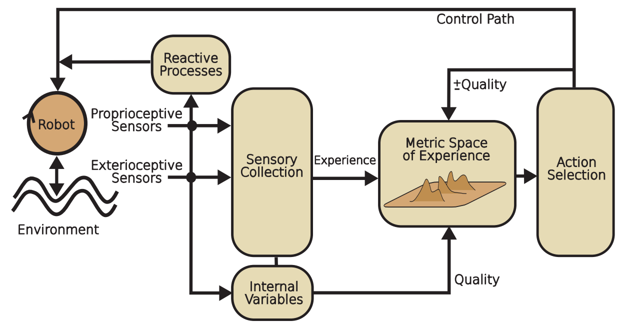

Figure 1 Interaction history based control architecture. See text for description.

There are two key aspects of this architecture. The first is the metric space of experience whereby new experiences appear as points in a growing and changing metric space. The second is the action–selection system. This closes the perception–action loop and also closes an internal loop feeding back and modifying the experience space. A quality measure, as determined by the agent's motivation and drives, is conferred onto each experience and that along with proximity in the metric space is used to distinguish experiences and select action. We describe these two aspects in the following sections.

3.1 Metric Space of Experience

Central to the proposed architecture is the capability to make metric comparisons between episodes of sensorimotor experience. An advantage of considering episodes is that they potentially hold more information about recent interactions than does the current sensorimotor state by itself.

One approach is to look for regularities in the statistical and informational structure of the data. Informational and statistical structure of sensorimotor data can also be used to characterize or "fingerprint" behavior (te Boekhorst, Lungarella, & Pfeifer, 2003; Tarapore, Lungarella, & Gomez, 2004) and also for a robot to classify its own behavior on-line using trajectories in sensor–motor spaces constructed from metric measures of distances between sensors (Kaplan

& Hafner, 2005; Mirza, Nehaniv, Dautenhahn, & te Boekhorst, 2005b).

In the following sections we describe the application of Shannon information theory (Shannon, 1948) to compare episodes of sensorimotor experience (see also Mirza, Nehaniv, Dautenhahn, & te Boekhorst, 2005a; Nehaniv, 2005). The basis is the information metric (Crutchfield, 1990), a measure of the distance, in terms of bits of Shannon information, between two information sources. We use the measure to compare sensorimotor experience over time and across modalities. Moreover, we close the loop to adaptive behavior by allowing the agent to act based on remembering its previous experiences in this space of its own temporally extended sensorimotor experiences. Here the notion of "temporally extended experience" is operationalized in a rigorous way using the flow of values over the agent's sensorimotor variables during a particular interval of time (temporal horizon).

3.1.1 Sensors as Information Sources An agent situated and acting in an environment will have many external and internal sensory inputs, any of which can be modeled as random variables changing over time. Consider one such random variable $\mathcal{X}$ changing with time, taking values $x(t) \in X$, where X is the set of its possible values. Time is taken to be discrete (i.e., t denotes a natural number) and $\mathcal{X}$ takes values in a finite set or "alphabet" $X = \{x_1, \ldots, x_m\}$ of possible values.[4]

Furthermore, any sensor or motor variable $\mathcal{X}$, beginning from a particular moment in time t_0 until a later moment $t_0 + h$ ($h > 0$), with the sequence of values $x(t_0), x(t_0 + 1), \ldots, x(t_0 + h - 1)$, can be considered as the time-series data from a new random variable $\mathcal{X}_{t_0, h}$, the sensorimotor variable with temporal horizon h starting at time t_0.

3.1.2 Information Distance For any pair of jointly distributed random variables (sensors) $\mathcal{X}$ and $\mathcal{Y}$, the conditional entropy $H(\mathcal{X}|\mathcal{Y})$ of $\mathcal{X}$ given $\mathcal{Y}$ is the amount of uncertainty that remains about the value $\mathcal{X}$ given that the value of $\mathcal{Y}$ is known, and is given by

$$H(\mathcal{X}|\mathcal{Y}) = -\sum_{x \in X}\sum_{y \in Y} p(x, y)\log_2\frac{p(x, y)}{p(y)},$$

where $p(x, y)$ is given by the joint distribution of $\mathcal{X}$ and $\mathcal{Y}$.

We assume approximate local stationarity of the joint distribution of random variables representing the sensorimotor variables over a temporal window and that this can be estimated closely enough by sampling the sensorimotor variables.

The information distance between $\mathcal{X}$ and $\mathcal{Y}$ is then given by

$$d(\mathcal{X}, \mathcal{Y}) = H(\mathcal{X}|\mathcal{Y}) + H(\mathcal{Y}|\mathcal{X}).$$

Crutchfield (1990) shows that this satisfies the mathematical axioms of equivalence, symmetry and the triangle inequality and so is a metric. Specifically, for three information sources $\mathcal{X}$, $\mathcal{Y}$ and $\mathcal{Z}$, d is a metric if it satisfies the following:

1. $d(\mathcal{X}, \mathcal{Y}) = 0$ if and only if $\mathcal{X}$ and $\mathcal{Y}$ are equivalent;
2. $d(\mathcal{X}, \mathcal{Y}) = d(\mathcal{Y}, \mathcal{X})$ (symmetry);
3. $d(\mathcal{X}, \mathcal{Y}) + d(\mathcal{Y}, \mathcal{Z}) \geq d(\mathcal{X}, \mathcal{Z})$ (triangle inequality).

Thus, d defines a geometric structure on any space of jointly distributed information sources.

Given two sensorimotor variables $\mathcal{X}_{t_0, h}$ and $\mathcal{Y}_{t_1, h}$ over a temporal horizon of window size h, we can estimate the information distance $d(\mathcal{X}_{t_0, h}, \mathcal{Y}_{t_1, h})$ by measuring the frequencies of occurrence of values $(x_{t_0 + i}, y_{t_1 + i})$ as i runs from 0 to $h - 1$.

With $t_0 = t_1$, $d(\mathcal{X}, \mathcal{Y})$ gives the information distance between different variables at the same time t. With $\mathcal{X}$ and $\mathcal{Y}$ taken from the same sensorimotor variable at different times, $d(\mathcal{X}, \mathcal{Y})$ gives the information distance between time-shifted regions of the variable.

Clearly there are issues related to the size of the temporal horizon h and also the number of values (bins) $\mathcal{X}$ and $\mathcal{Y}$ may take that affect the accuracy of these estimates. These issues are examined in Mirza et al. (2005a) showing that behavior can be categorized robustly over a wide range of numbers of bins and horizon lengths.

3.1.3 Experience and the Experience Metric Given the above definitions we can now formalize an agent's experience from time t over a temporal horizon h as $E(t, h) = (\mathcal{X}^1_{t, h}, \ldots, \mathcal{X}^N_{t, h})$ where $\mathcal{X}^1, \ldots, \mathcal{X}^N$ is the set of all sensorimotor variables available to the agent.

We can then define a metric on experiences of temporal horizon h as

$$D(E, E') = \sum_{1=k}^{N} d(\mathcal{X}^k_{t,h}, \mathcal{X}^k_{t',h}),$$

where $E = E(t, h)$ and $E' = E(t', h)$ are experiences of an agent at time t and t' over horizon h, and d is the information distance. The fact that D is a metric follows from the fact that the metric axioms hold component-wise, as d is a metric.

As experiences are collected, they can be placed in a metric space of experience using the experience metric. The maximum dimensionality of the space is $N - 1$, where N is the number of experiences in the space.

3.2 Action Selection

A simple mechanism is adopted for action selection whereby the robot can execute one of a number of atomic actions (or no action) at any time step. This is seen as a tractable first step, and a more sophisticated action or behavior generation capability would allow for more open-ended development.

The actual action selected will be either a random selection of one of the atomic actions or an action that was previously executed after an experience in the history that is near to the current episode. An advantage of this approach is that behavior can be bootstrapped from early random activity, and later behavior built on previous experience.

The process of action selection is as follows.

1. Up to K candidate experiences from the experience space within a given information distance radius[5] r_0 of the current experience $E_{current}$ are initially selected.
2. These K experiences are ranked as $E_1, \ldots, E_K$ according to how close they are to $E_{current}$.
3. Then, sequentially, experience E_i is chosen with probability a linear function of the quality of E_i until either an experience is chosen or the ranked list is exhausted.
4. If an experience is chosen from the candidate list, then the particular action that was executed following the chosen experience is then chosen as the action to be executed next; otherwise a random action is chosen.

The exact nature of the calculation of quality is dependent on the nature of the intrinsic drives and motivations ascribed to the agent. For the experimental scenarios used in this article, a specific motivational system was designed (see Appendix A). However, we note that this could be altered and generalized for other types of interaction.

The linear mapping from quality to probability ensures that, with small probability, the robot may still choose a random action, as this may potentially help to discover new, more salient experiences. This has the advantage of emulating body-babbling, that is, apparently random body movements that have the (hypothesized) purpose of learning the capabilities of the body in an environment (Meltzoff & Moore, 1997). Early in development, there are fewer, more widely spread experiences in the space, so random actions would be chosen more often. Later in development, it is more likely that an the action selected will come from past experience.

Finally, a feedback process evaluates the result of any action taken in terms of whether there was an increase in quality after the action was executed. It then adjusts the quality of the candidate experience, from which the action was derived, up or down accordingly. Using this mechanism, the metric space is effectively altered from the point of view of the action-selection system. Closing of the perception–action loop in this way with feedback, together with growth of the experiential metric space, results in the construction of modified behavior patterns over time. This can be viewed as a form of ontogenetic development and adaptation, that is a process of change in structure and skills through embodied, structurally coupled interaction.

3.3 Implementation

The interaction history architecture was implemented using using URBI (Baillie, 2005) and Java on a Sony Aibo ERS-7 robot dog and a personal computer running the Linux operating system. URBI provides the robot control layer and a full-featured event-based parallel scripting system. The URBI software runs directly on the robot where actions and background

Table 1 Sensors and internal variables.

Type	Examples	Total
Exterioceptive	IR-distance, buttons	15
Visual	Average color values in a 3×3 grid over image	9
Proprioceptive	Joint positions	18
Internal	Face position, ball position, desire to see a face	10

behaviors are executed; URBI receives and processes events and controls motors every 35 ms. The system runs on-line with telemetry data and video images being sent over wireless to the personal computer approximately every 80–120 ms, where the metric space of experience is constructed and used in action selection. We define a time step upon reception of each set of data, so the time between time steps varies and is approximately 80–120 ms.

The sensory information available to the robot falls into three broad categories: proprioceptive (from motor positions), exterioceptive (environmental sensors, including vision) and internal (these might, for instance, indicate drives and motivations, or be the result of processing of raw sensory data, e.g., ball position). Vision sensors are built by subdividing the visual field into regions and taking average color values over each region at each time step. In this implementation a 3×3 grid over the image is used, taking the average of the red channel only, resulting in nine sensors for vision. A generalized human face detection system, required for the interaction experiments of Section 6 was implemented using Intel OpenCV HAAR Cascades (OpenCV, 2000), smoothed to remove short gaps (< 50 ms) in detection. The variables used in this implementation are summarized in Table 1, with further discussion of internal variables in Appendix A. Note that audio is not used in these experiments.

The basic object of data in the architecture is an experience. For every experience, the quantized values of all sensors over the time horizon h are required to determine the information distance between the experience and any other one, and so are stored. Additionally, the quality value of the experience, as determined by the motivational system detailed in Appendix A, is stored with each experience, subject to modification in interaction as described in Section 3.2.

The horizon length h of the experiences used to construct the metric space and the number of bins Q used to quantize sensor data are parameters set for each particular experiment. Experiences are taken from the sensorimotor data stream every G time steps, where G is the experience granularity. Thus, a granularity of $G = 2$ would store an experience of h time steps at every other time step.

The metric space is continually being updated as new experiences are added, by calculating the experience distance between the new experience and all other experiences in the space. For efficiency, a list of near experiences is kept for each experience and is updated as new experiences are added.

A list of actions being executed (if any) at any time step is kept and consulted when determining what actions were executed immediately following any given experience.

4 Related Work

There are many potential architectures that take history of action and interaction into account. Top-down deliberative architectures, such as ACT-R, include memory storage and retrieval and others, such as Soar, have been extended to include episodic memory (Nuxoll & Laird, 2004). In the model of Nuxoll and Laird, the features of the episode are encoded and used in retrieval by matching. This external representation of sensory input is common. Connectionist systems that have memory include, for instance, Elman networks or recurrent neural networks. Rylatt and Czarnecki (2000) have shown that generally recurrent neural networks are not well suited to learning delayed response tasks. Additionally, recurrent networks are very hard to design beyond a certain size, and this requires that sensory input be encoded and reduced in quantity.

Approaches such as echo state networks and liquid state machines attempt to address this limitation by training only the output nodes of a network (Jaeger & Haas, 2004). The memories of episodes appear only as weights and attractors of the system and so different episodes cannot be compared. Other approaches include certain behavior-oriented control systems combined with learning (Matarić, 1992; Michaud & Matarić, 1998). Most behavior-based models do not include learning from past experience, but of those that do our approach differs in that the history is not specified in terms of the behavior being selected (or indeed, the action being selected), but in terms of the sensorimotor history.

Our work is related to reinforcement learning (Sutton & Barto, 1998), particularly those examples that use intrinsic motivation (e.g., Barto & Şimşek, 2005; Bonarini, Lazaric, Restelli, & Vitali, 2006). Our approach, however, uses temporally extended experience rather than the instantaneous values of the sensorimotor and internal variables (state). We would argue that this distinction is important, as temporal structure is inherently captured in experiences of different lengths. Moreover, we do not assume that the environment can be modeled as a Markov decision process (MDP; this is particularly important when there is an interaction partner) as is the case with most reinforcement learning paradigms and, in particular, with approaches that do not use a model (e.g., Q-learning). Furthermore, our approach does not require a static state space to be circumscribed at the outset, but instead uses a growing and changing space of experiences, where potentially in the course of ontogeny the set and character of sensors, actuators, and embodiment may change.

Related work in the multi-agent domain (Arai, Sycara, & Payne, 2000) has agents in a grid world acquiring coordination strategies, and uses a fixed-length episodic history expressly to counter the MDP assumption. However, this model is also state-based and so uses a profit-sharing mechanism to assign credit to state–action pairs. Moreover, it does not compare episodes of history with previous ones, or locate them in a metric space.

Examples of a developmental approach used in robots include Blank, Kumar, Meeden, and Marshall (2005), where a robot uses subgoals to develop smooth-wall following in an architecture that uses self-organizing maps of visual and sonar data, and Oudeyer, Kaplan, Hafner, and Whyte (2005), where an Aibo robot discovers object affordances through an "adaptive curiosity" driven developmental framework. Kaplan and Oudeyer (2006) also propose mechanisms of drive and motivation based on "progress niches," which allow an agent to maximize learning and developmental progress in a way analogous to Vygotsky's "zone of proximal development" (Vygotsky, 1978).

As interest in developmental robotics gains momentum, we will increasingly see play-like scenarios used to scaffold early development of robots (Oudeyer et al., 2005), to study human cognitive development (Kozima, Nakagawa, & Yano, 2005) and just for entertainment (Brooks et al., 2004). Likewise, our use of an interaction game, peekaboo (see Section 6.1), played by human children during early development was deliberately chosen to bring robotic development closer to human development. See also Dautenhahn, Bond, Cañamero, and Edmonds (2002) for a representative review of robots socially interacting in play.

Recent research has used information methods in the analysis and control of (simulated and unsimulated) robot behavior. Lungarella and Sporns (2005) use informational measures (including mutual information and a related complexity measure) to quantify the degree of statistical structure in sensorimotor spaces, and suggest that perceptually guided movement generates high degrees of regularity and correlation. Olsson, Nehaniv, and Polani (2004) use an information distance measure to find structure in uninterpreted sensorimotor data, and also show that this is superior to other measures such as the Hamming metric and the correlation coefficient (Olsson, Nehaniv, & Polani, 2006b). In particular, they show that information measures are a general method for quantifying functional relationships between sensorimotor variables, including non-linear relationships, which we note may be important in systems situated in complex, real environments. Having learnt how its sensorimotor system is structured through information self-structuring during coordinated sensor–motor action, it is possible for a robot to learn how its effectors can be used, for example, for simple motion tracking (Olsson, Nehaniv, & Polani, 2005, 2006a). In earlier work, Pierce and Kuipers (1997) achieve learning of sensory maps and motor control laws from uninterpreted sensors and effectors by use of statistical structure in the data rather than informational methods.

5 Experimental Validation of Metric Space of Experience

In this first experiment, the metric space of experience was tested in the absence of the action control loop (although experiments in the next section include this loop). For the metric space to be useful in an interaction history, experiences that appear to be similar by a suitable subjective measure must also be close according to the measure of distance used to place experiences in the interaction histories metric space. To test this, the history is used to predict the future path of a ball based on recent sensory experience. If the experiences are well matched, then so will the predicted path.

5.1 Validation Experiment: Experimental Setup

The robot was stationary in a "sitting" position, with the head pointed forward. A pink ball was moved in the air in view of the robot's head camera at a distance of approximately 30 cm. The path of the ball in each trial included repeated vertical, horizontal and circular movements.

The sensory data collected included the horizontal and vertical locations of the ball with respect to the video frame (calculated using simple color thresholding) along with the full sensorimotor input of Table 1. In addition, the ball position at the end of each episode of experience was stored along with each experience. The predicted future position of the ball was then taken from the positions stored with the experiences following the nearest previous experience to the current one.

It is important to note that the robot is not matching current ball position with previous ball position, rather we use all sensory and motor variables as information sources to detect similarity between experiences, and then use the stored ball position to give the experimenter an indication as to how well the experience was chosen. For verification purposes, a path is drawn on the display of the robot's visual field during operation, indicating the predicted future path.

The horizon length of the experiences was 40 time steps or approximately 3,400 ms (a single time step was approximately 85 ms long). The data were quantized into five bins in the probability distribution estimation algorithm. The ball was moved such that the time for the ball to describe a circle (or to move horizontally or vertically for a complete cycle) was 6–7 s. Thus the horizon length was shorter than, but on the same scale as, a single cycle of the repeated behavior, and the experiences would comprise approximately a half of a cycle.

5.2 Validation Experiment: Results and Analysis

Figure 2 shows a sequence of images from one trial with one image shown per experience. The sequence lasts just over 4 s and consists of approximately 50 time steps (one time step ~85 ms) and 12 experiences (experience granularity G of four time steps). There were 112 overlapping experiences (about 39 s of activity) before those shown, during which the ball was moved from left to right four times and in a circle once. Each image shows the robot's camera view during an experience with the predicted path overlayed (at run-time). For clarity, a single image from the sequence is reproduced in Figure 3 with the position of the ball and the predicted path highlighted.

In the sequence shown and others, the robot required very few examples of a sequence (usually one) before the appropriately predictive experience could be located. This demonstrates that the information distance measure is capable of placing subjectively similar experiences (to an external observer) near to each other in the experience space (of the agent). However, it was found that while the path of the ball could be predicted fairly well early on in the sequence, later on, as the choice of experiences grew, the candidate experience chosen was not always the most appropriate.

Occasionally, subjectively inappropriate experiences were matched. As an example, consider the seventh image in Figure 2. Here, the predicted path inferred from the sequence of experiences following the candidate experience corresponds to the half-circle that the ball has just been through (rather than the half-circle it is just about to go through, as in the other images). The candidate experience chosen is informationally close to another experience half a cycle back in time, which may have been more appropriate. These two possible experiences that could have been matched correspond to motions of the ball from opposite sides of a circle. As the experience distance measure is the sum of information distances between

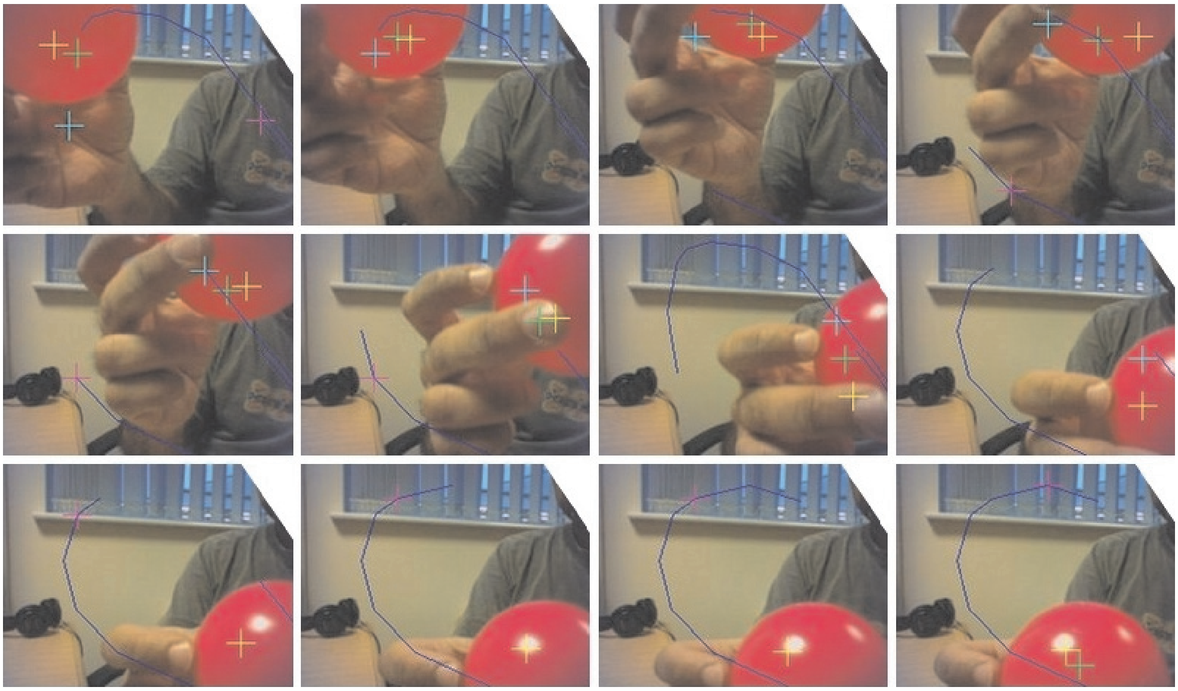

Figure 2 Validation experiment. A series of 12 consecutive images from the Aibo camera showing ball path prediction using a sensorimotor interaction history. The robot does not move its head in this sequence. Images are sequential left to right and top to bottom. The sequence lasts ≈ 4.2 s (49 time steps or 12 experiences) and is taken after 38 s of activity. The line shows the path prediction for 10 experiences ahead. The crosses are from various methods for ball detection; only one of these was actually used as sensory input. Horizon = 40, number of bins = 5, experience granularity = four time steps. One image shown per experience.

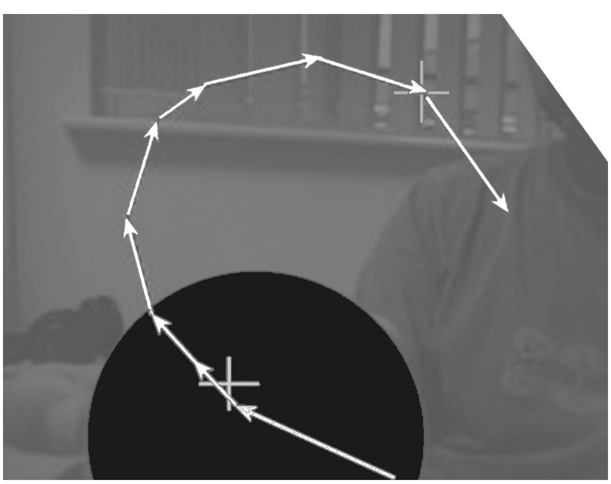

Figure 3 Single image from the Aibo camera taken during the ball prediction experiment. The predicted path has been highlighted with arrows, starting from the position of the ball during the matched experience, and ending with the position of the ball during the tenth experience after the matched one. The lower cross-hair is the detected ball position, and the upper cross-hair is the predicted ball position.

variables, then a symmetric error such as this is likely, especially as phase-shifted periodic variables can have a small or zero[6] information distance. This particular test scenario presents an unrealistic situation where the robot does not move, and we predict that with embodied action, more information would be available with which to distinguish experience.

6 Interaction Game Experiments

In this section we describe two experiments that use the experience metric space in a robot that develops the capability to play a simple interaction game. In the first, a human partner engages in a peekaboo game with a robot, and in the second the effect of the experience horizon length on the ability of a robot to develop the capability to play the game is investigated. We describe and motivate the choice of the peekaboo game as an interaction scenario for this study, followed by a description of the experiments and results.

6.1 Peekaboo Early Interaction Game

The development of gestural communicative interaction skills is grounded in the early interaction games that infants play. In the study of the ontogeny of social interaction, gestural communication and turn-taking in artificial agents, it is instructive to look at the types of interaction that children are capable of in early development and how they learn to interact appropriately with adults and other children. A well-known interaction game is peekaboo where, classically, the caregiver, having established mutual engagement through eye-contact, hides their face momentarily. On revealing their face again, the caregiver cries peek-a-boo!, peep-bo!, or something similar. This usually results in pleasure for the infant, which, in early development, may be a result of the relief[7] in the return of something considered lost (i.e., the emotionally satisfying mutual contact), but later in development also may be a result of the meeting of an expectation (i.e., the contact returning as expected along with the pleasurable and familiar sound), and the recognition of the pleasurable game ensuing (Montague & Walker-Andrews, 2001; Veatch, 1998).

Bruner and Sherwood (1975) studied peekaboo from the point of view of play and learning of the rules and structures of games. They also recognized that the game relies on (and is often contingent with) developing a mastery of object permanence as well as being able to predict the future location of the reappearing face. We suggest that the parts of the game can be viewed as gestures in a non-verbal communicative interaction. The hiding of the face is one such gesture; the vocalization, and the showing of pleasure (laughing) are others. In order for the interaction game to proceed successfully, the gestures must be made by either party at the times expected by the players, and that absence or mis-timing can result in the game cycle being broken. Learning of the game is supported by further gestures, such as a rising expectant intonation of the voice during hiding, as a reassurance or cue of the returning contact. Later in development, the roles of the game can become reversed with the child initiating the hiding, while still obeying the established rules by, for instance, uttering the vocalization on renewed contact.

In all this, the rhythm and timing of the interaction are crucial and Bruner and Sherwood have suggested that the peekaboo game and other early interaction games act as scaffolding on which later forms of interaction, particularly language and the required intricate timing details, can be built (Pea, 2004, pp. 424–5).

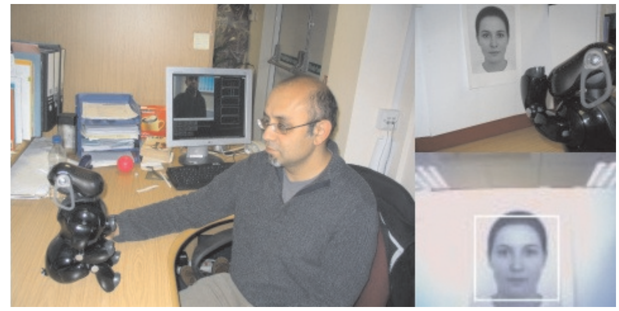

Figure 4 Aibo playing the peekaboo game. Left: Sony Aibo with human partner. Right: Using a static image. (Top: hiding head with front leg. Bottom: Aibo's view, showing face detection.)

In relation to the development of social cognition in infants, "peekaboo" and other social interaction games, which are characterized by a building and then releasing of tension in cyclic phases, are important as they are considered to contribute developmentally to infant understanding and practise of social interaction. Peekaboo provides the caregiver with the scaffolding upon which infants can co-regulate their emotional expressions with others, build social expectations and establish primary intersubjectivity (Rochat, Querido, & Striano, 1999).

6.2 Interaction Experiment 1: Sensorimotor Contingencies in the Interaction Game Peekaboo

The purpose of this experiment was to investigate whether an embodied interaction history in a robot could be used for the robot to act appropriately in an interaction that requires following a spatio-temporally structured set of rules, which when followed result in a high value according to an internal motivational system.

6.2.1 Interaction Experiment 1: Experimental Setup The robot remains in a sitting position (see Figure 4) throughout the experiment with the forelegs free to move, facing the human interaction partner at a distance of 30–50 cm. The actions that the robot can

Table 2 Actions.

Action	Description
0	Do nothing
1, 2	Look right/left
3	Track ball with head
4, 5	Re-center head
6, 7	Hide head with left/right foreleg
8, 9	Wave with left/right foreleg
10	Wag tail

execute are listed in Table 2. Each action takes 2 s or less and the re-center head action is duplicated to offset the two actions that take the head away from the center.

The human partner takes a passive role with the usual interaction feedback from the partner provided by an internally generated motivational value in the robot. The action to hide head with foreleg means that the robot covers its forward-facing camera with one or other of its forelegs, before uncovering it again a short time later.

In this experiment and the next, we define a "peekaboo sequence" to have occurred when the robot having detected a face, through action loses detection and returns to detect the face again, with this cycle repeating at least once. This is marked, because of the nature of the motivational dynamics (see Appendix A), with a high value for the motivational variable m. The duration of the sequence is measured from the point of the first loss of face detection through to the last point at which high motivation can be sustained without a break in the sequence. The average cycle period is the average duration of a single face loss/re-detection cycle within a peekaboo sequence.

6.2.2 Interaction Experiment 1: Results and Analysis

There were 15 trials conducted, each lasting between 3 and 5 min. The results tend to show that the robot, after a period of random movement, does start to engage in repeated cycles of behavior. In 10 of the trials, the robot engages in peekaboo as defined above. If the robot were not to take action to block its own camera view, it would have long periods of detecting a face, which does not result in a high value for the motivational variable. Instead, the robot generates intermittency in detecting a face by executing actions 1, 2, 6 or 7 in Table 2. Figure 5 shows the trace of the internal variables as well as the actions executed from one short trial where peekaboo behavior was observed. The sequence consists of eight repeated cycles of hiding interspersed with other actions, which importantly include actions to re-center the head.

The trials also show that it is easy for the robot to "get stuck" in areas of the experience space, especially if all other factors in the environment remain unchanged. This occurs four times in these trials, usually with the robot repeating an action such as waving.

The results also show that relatively few experiences are selected and thus modified (with regard to their stored quality value) over time. In some of the trials, particular experiences were selected multiple times, but this is not always the case. In the trial of Figure 5, 34 choices of action were made: the first 11 were random actions, and 13 of the remaining 23 actions were selected from a total of 12 previous experiences (the other 10 being randomly selected).

6.3 Interaction Experiment 2: Investigation of the Effect of Horizon Length

First, the purpose of this investigation was initially to evaluate whether the model for development based on interaction history performed better than random for the task of playing the game of peekaboo. Second, the hypothesis that the horizon length of experience would affect the ability to acquire peekaboo behavior was tested by trying a number of different horizon lengths in a controlled experiment. The hypothesis was that the horizon length of experience needs to be of a similar scale to that of the interaction in question. If it is too short, the experience does not carry enough information to make useful comparisons to the history. If it is too long, then the interesting part of the interaction becomes lost in the larger experience.

6.3.1 Interaction Experiment 2: Experimental Setup

Again the robot remains in a "sitting" position throughout the experiments but facing instead a picture of a face (see Figure 4) at a fixed distance of

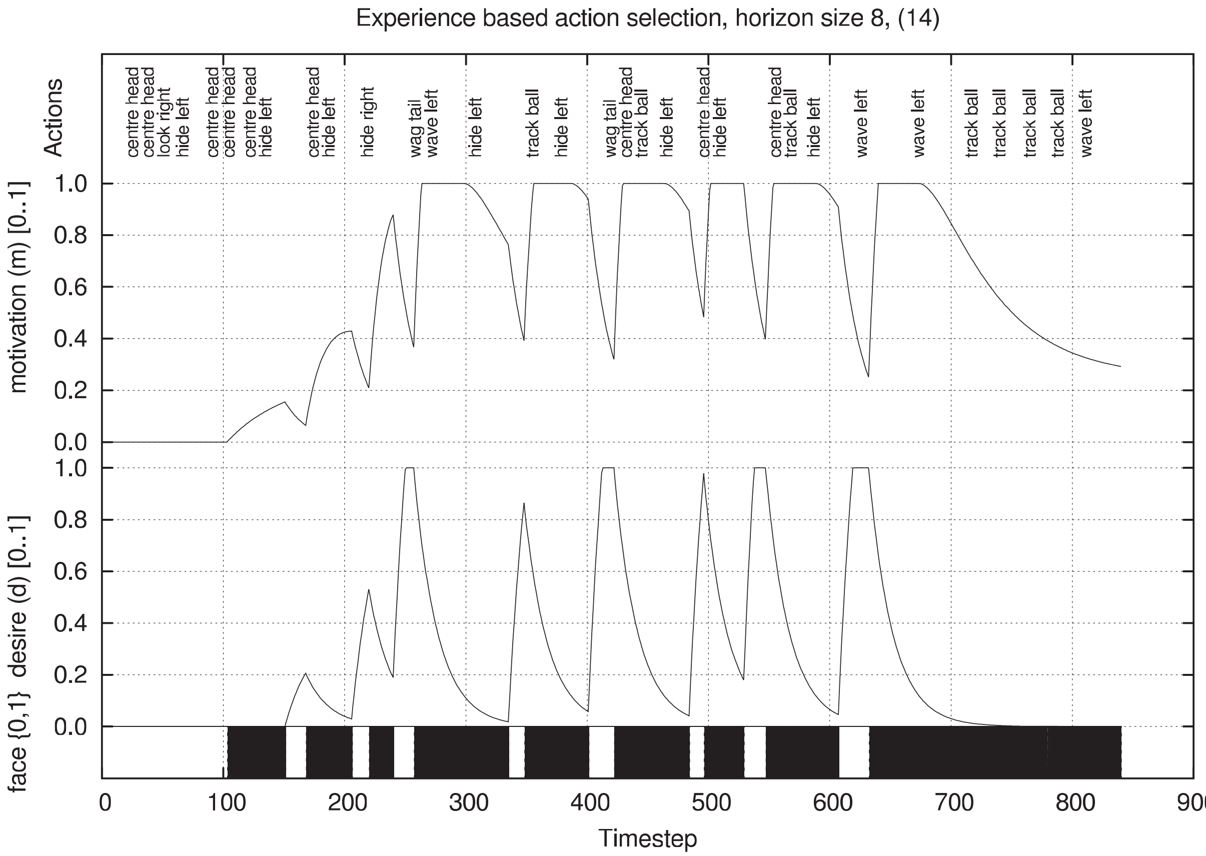

Figure 5 Time series of motor and sensor values showing the engagement of the robot in the peekaboo game. The bottom part of the graph shows when the face is seen, and the two internal variables are shown varying in response to this. The actions executed are shown at the top of the trace.

40 cm. A picture was used rather than an interaction partner in these particular experiments to allow analysis of the robot's interactions in isolation when comparing horizon lengths, and for experimental repeatability.

We ran six trials of 2-min duration for each horizon length of 8, 16, 32, 64, and 128 time steps (0.96, 1.92, 3.84, 7.68, and 15.36 s, respectively). For comparison, a further six trials were run where the choice of action was random and not based on history. In each of the trials, the metric space started unpopulated.

6.3.2 Interaction Experiment 2: Results Table 3 summarizes the results of 36 trial runs, while Figure 6 shows, for selected trials, time-series graphs of the motivational variables coupled with the actions taken. Peekaboo behavior, as defined in Section 6.2.1, was seen in 18 of the 36 runs. All but one of the horizon size 8 trials, and four of horizon size 16, also showed peekaboo behavior. The sequences were mostly generated by repetitive actions for long durations. Figure 6A (horizon size 8) shows the best example of this behavior; the average cycle period is approximately 42 time steps or 5 s, and the sequence duration is around 640 time steps (76 s). During this sequence, the head is hidden to the left and right, and this is interspersed with head-centering actions. Through all of these episodes, periods of no action serve to alter the timing of the cyclic periods. Although all of the trials using random action selection showed some peekaboo behavior, they were irregular both in terms of cycle period length and in terms of the actions used to generate the sequence (see Figure 6B for example).

Of the longer horizon length (32, 64, and 128) trials, three showed peekaboo behavior using repeated

Table 3 Experiment summary. Duration and average cycle period in time steps (ts) of peekaboo sequences for each trial. Where peekaboo is achieved using a waving instead of hiding action, this is indicated as waving.

Run	Random length/period (ts/ts)	Horizon 8 length/period (ts/ts)	Horizon 16 length/period (ts/ts)	Horizon 32 length/period (ts/ts)	Horizon 64 length/period (ts/ts)	Horizon 128 length/period (ts/ts)
1	120/40	180/45	260/40	None	400/57 waving	None
2	220/55	150/40	None	None	None	None
3	220/45	*Fig 6A* 640/42	140/45 200/50	*Fig 6F* None	None	100/40
4	200/60	130/45 150/70	*Fig 6E* 260,240/40	None	None	None
5	160/50	None	140/35 waving	*Fig 6C* 540/47 waving	*Fig 6D* 220,100/37 100/40	120/40
6	*Fig 6B* 80,140/40	250/42	120/40	840/47 waving	None	None

actions (e.g., Figure 6D). Three also showed peekaboo using an action (waving) that would not normally cause a break in face detection. In this particular circumstance, "rocking" of the robot caused a break in face detection > 50 ms and led to a peekaboo sequence (see Figure 6C for an example.)

6.3.3 Interaction Experiment 2: Analysis All of the trial runs of random action selection resulted in some peekaboo sequences, although with mixed, irregular actions. It is likely that this is due to a motivational system that responds to a wide range of frequencies[8] combined with a range of actions, four of which would result in some loss of face detection. However, to see longer peekaboo sequences with regular actions, some controlled behavior must be selected and this is only seen in the experience-driven trials. As a contrary example, see Figure 6F where no peekaboo-like dynamics are seen.

In some of the experience-driven trials, repeated behavior was seen that could have resulted in high motivation if the head had been pointed forward. Experience alone was not able to re-center the head. On one occasion however, when the head was re-centered (randomly), then the experience space allowed a resumption of the peekaboo sequence (see Figure 6E). Thereafter, a re-centering action is selected along with hiding actions.

Figure 6 Motivational dynamics and actions for selected 2-min interaction sequences of different horizon lengths. Graphs show when face is seen (black bars at bottom), the values of the key internal variables (m and d), and the action taken at the top (note that action 0, do nothing, is not shown for clarity). (A) Peekaboo. Horizon size 8. Dynamics during an extended peekaboo sequence. (B) Random action selection resulting in high m and d. Although the action selection is random, it is possible to obtain periods of high value. (C) Emergent behavior resulting in high m and d. Horizon size 32. Dynamics generate a high value when the face is intermittently lost when the waving paw returns to hit the hind knee and jogs the robot. (D) Irregular response to regular actions. Horizon size 64. The regular hiding of the head does not always result in a high value, perhaps because the face is not detected during the period that the head points forward. (E) Repeated sequence. Horizon size 16. Sequence of peekaboo repeated after the head is re-centered. (F) Peekaboo not inevitable. Horizon size 32. Here, although the head is hidden twice, the peekaboo dynamics are not inevitable and coordinated action is necessary for continued high motivation.

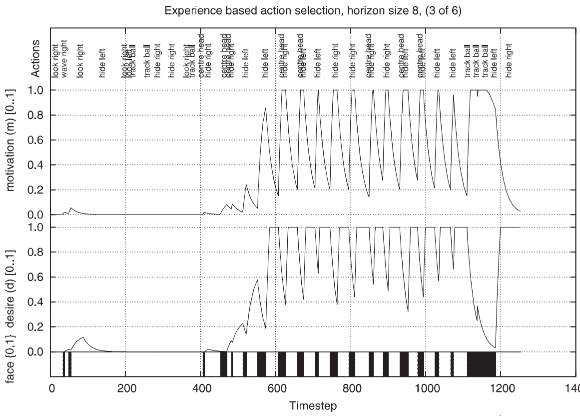

A: Horizon 8, run no. 3/6

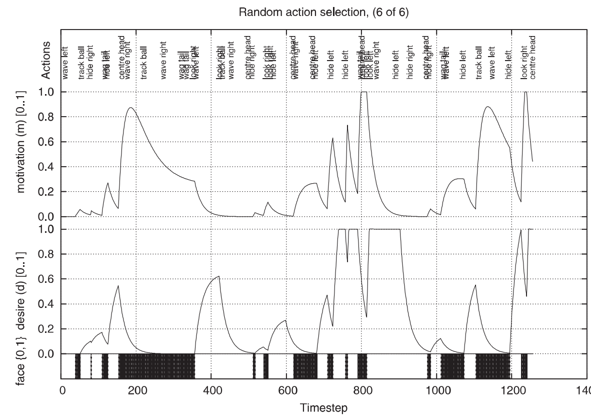

B: Random, run no. 6/6

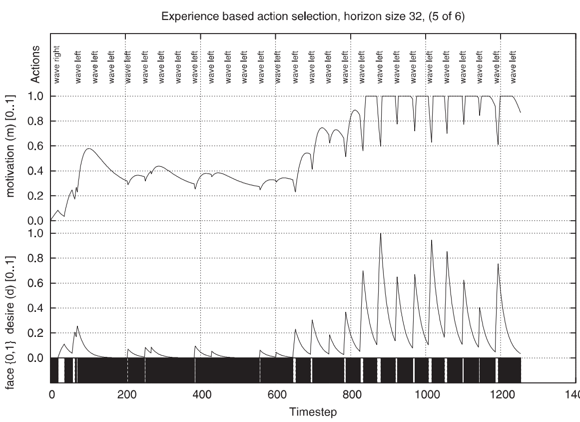

C: Horizon 32, run no. 5/6

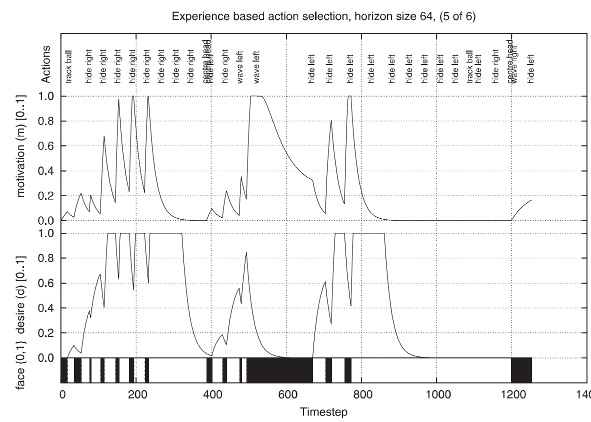

D: Horizon 64, run no. 5/6

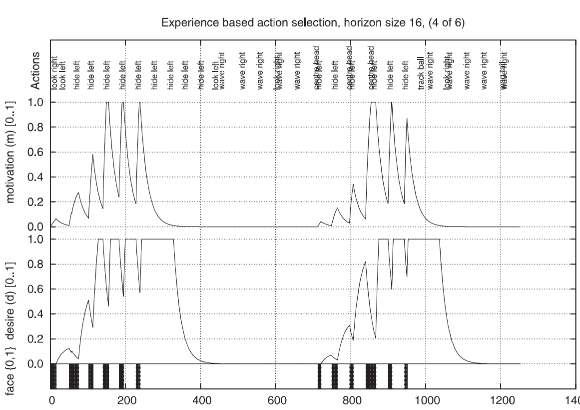

E: Horizon 16, run no. 4/6

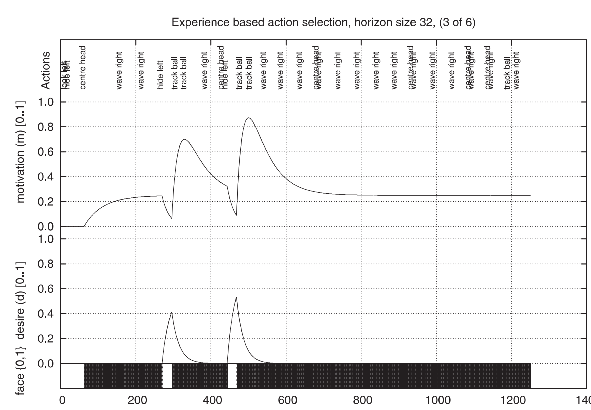

F: Horizon 32, run no. 3/6

The best of the cyclic behavior was seen in the experience-driven trials of horizon size 8 and 16 time steps (about 1 and about 2 s, respectively). This result indicates that it may be necessary to have an appropriately sized time-horizon, and this may be related to the length of single actions (about 2 s), and thus the natural period[9] of the cyclic behavior. This could be because to bootstrap the initial repetitive behavior, it is necessary to focus on an experience of one cycle length when there is only a single (possibly randomly generated) example of the cycle in the agent's experience.

7 Summary

We have motivated and presented a definition of grounded sensorimotor interaction histories for embodied organisms, and presented a control architecture for an artificial organism using such a history. We also argue that a system that connects action with dynamically constructed experiences can scaffold ontogenetic development, given a sufficiently sophisticated system of goals and motivations.

The system was implemented in an Aibo robot, and results from a validation experiment have confirmed that a metric space of experience based on information distance measures between time-extended episodes of sensorimotor experience might be a suitable basis for extending the temporal horizon using interaction histories in robots. Experiments using a robot playing a simple interaction game using this architecture have shown that it is able to develop the capability to play the game based on its own experience and an internal motivational system. Further results indicate that the horizon length of experience plays an important role in the types of interaction that can be engaged in. The experimental results support the hypothesis that horizon length needs to be of a similar scale to that of the interaction in question, and thus should be determined, at least in part, by the types of interaction that will take place. The action-selection architecture is however still extremely limited, and this combined with the short experiment lengths and the oversensitive motivational system suggests various directions for improvement.

8 Future Work

An important direction that needs to be explored is the anticipation of future action and expectation of future reward, although how far ahead in the future may vary for the development of different skills and task abilities. Currently, experiences of the same length are being compared; however, it is also possible to have shorter-term current experience being matched against parts of longer-term episodic experience, and the current short experience being assessed with an anticipated future value related to the best value in the extended experience. We expect this approach to better balance the requirement, as found above, to have horizons of appropriate size for comparing experience successfully, while also taking into account temporally extended aspects of interaction.

Further, given the apparent dependence on horizon length, it may be necessary to operate on many different horizon lengths, and an adaptive, variable experience length may help in then finding areas of high value for the different types of interaction the robot will encounter. We suggest that an approach to deciding on appropriate experience lengths will come from the density of "interesting" features or events in the experience space, the operational determination of which will take into account motivational dynamics, value of experience, and possibly rates of change of experience distances.

These particular experiments carried out so far do not have much non-trivial interaction with either the environment or the partner's side, and lack features of more contingent social interaction. However, their purpose was to establish the feasibility of using temporally extended experience-based interaction history architecture in adaptive behavior in controlled studies. The next steps must be to increase the social complexity of interactions using the interaction history approach (most likely requiring a more sophisticated motivational system) in less controlled scenarios, and to demonstrate further capacity for scaffolding the ontogeny of interaction skills in the social environment.

The current architecture is expensive in terms of both computer memory usage (increasing linearly with time) and computational complexity (increasing quadratically with time), and this cannot be sustainable in support for long-term development. A solution may be to reduce the number of experiences by "forgetting" (i.e., removing "unused" experiences from the metric space over time) or by "merging" similar experiences.[10] If a constant number of experiences were retained, then both memory and computational complexity would remain constant. Questions arise as to how many experiences to retain, and which to remove. It would be essential, however, to retain a sense of the structure of the experience space, and in particular the local density of experience.

We expect the structure of the dynamically growing and changing experience space to reveal important information about familiarity of experience, novelty of experience, areas of high and low reinforcement, areas of mastery and zones where current development can proceed through learning. Moreover, from the structure of the experience space, natural representations may emerge grounded in an agent's sensorimotor history developed through interaction, which are useful for ongoing developmental progress. Indeed, as areas of familiarity, mastery and novelty are identified, these may themselves provide a more general intrinsic motivational system that can drive development.

Acknowledgments

This work was conducted within the European Union 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.

Appendix A: Motivational Dynamics

We present the dynamic system of coupled equations that describe the motivational system used to confer a quality measure to experience. This feedback from the environment was designed specifically for the requirements of a peekaboo game, but could be generalized to other types of interaction.

To provide appropriate feedback, we require a high value for motivation following a period of peekaboo-like interaction. This is achieved by the interplay between a signal originating in the environmental interaction (perception of a face) and two internal variables.

First, the agent possesses a binary meta-sensor f that is a result of processing the visual sensors (image) to locate a generalized human face shape in the image, if one exists. Face detection is implemented using Intel OpenCV HAAR Cascades (OpenCV, 2000). This is then smoothed to remove short gaps (< 50 ms).

Second, the desire to see a face is given by d (constrained in the range [0, 1]) and increases when there is no face seen at a rate determined by how often a face has been seen recently (actually by feedback from m described below). The desire decays otherwise. See Equation 1.

Finally, the overall motivation m, also constrained in the range [0, 1], increases when $f = 1$. The rate of increase is determined by the desire to see a face d. In the absence of desire d, when a face is seen m tends to a constant value set by $C_{\max}$. When no face is seen, m decays at rate δ_3. See Equation 2.

In the experiments described in this article m is used as the quality value for the experiences.

$$\Delta d = \begin{cases} \alpha_1 m - \delta_1(1-m) & \text{if } f = 0, \\ -\delta_2 d & \text{if } f = 1. \end{cases} \tag{1}$$

$$\Delta m = \begin{cases} -\delta_3 m & \text{if } f = 0, \\ \alpha_2 d + \beta(C_{\max} - m) & \text{if } f = 1. \end{cases} \tag{2}$$

$$d, m \text{ constrained such that } d, m \in [0, 1]$$

The parameters of the dynamics equations are shown in Table A1 along with the values used in the experiments. These values were chosen by trial and error, and we note that with these values the system is receptive to a wide range of periods for peekaboo.

Notes

1 This is true for the simpler vehicles that do not have a memory.

2 For instance, the bacteria *Escherichia coli* are known to have a certain minimal level of embodiment (Quick, Dautenhahn, Nehaniv, & Roberts, 1999) and "cognition" (van Duijn, Keijzer, & Franken, 2006), and are able, without a nervous system, to exploit fairly simple sensor–motor coupling through limited low-bandwidth channels to achieve reactive behavior such as chemotaxis.

3 The examples here are chosen from the separate but related fields of psychology, cognitive science, and artificial intelligence.

4 The approach generalizes to continuous time and value sets with appropriate changes.

Table A1 Parameters of dynamic equations for motivational system.

Parameter	Description	Value
α_1	Rate of increase of d based on m	0.12
α_2	Rate of increase of m based on d	0.12
C_{max}	Value that m tends to after long periods of $f = 1$	0.25
β	Rate that m tends to C_{max}	0.02
δ_1	Rate of decay of d when no face is seen	0.05
δ_2	Rate of decay of d when a face is seen	0.05
δ_3	Rate of decay of m when no face is seen	0.05

5 In these experiments the radius is fixed, but we note that this could be adapted on-line.

6 Variables that have a zero information distance are recoding equivalent and are not necessarily identical (see Crutchfield, 1990).

7 In the context of humor, peekaboo in its early stages is an example of relief laughter. That is, relief that the caregiver that is thought to have disappeared actually has not (Veatch, 1998).

8 The motivational system tuned with the parameters given in Appendix A would result in high values of the variable m after a few cycles where the face signal was lost for anywhere between 50 ms to 9.5 s. Thus, it was inevitable that a high motivational value should be reached with even random actions.

9 Note that the motivational system itself does not dictate this period as any cyclic behavior of period up to 19 s can result in high values of m.

10 Alternatives are to store fewer experiences in the first place and to make fewer comparisons, maybe assimilating and deleting some of these experiences during a "sleeping" phase.

References

Arai, S., Sycara, K., & Payne, T. R. (2000). Experience-based reinforcement learning to acquire effective behavior in a multi-agent domain. In *Proceedings of the 6th Pacific Rim International Conference on Artificial Intelligence*, Melbourne, Australia, (pp. 125–135). *Lecture Notes in Artificial Intelligence* Vol. 1886. Berlin: Springer.

Baillie, J.-C. (2005). URBI: Towards a universal robotic low-level programming language. In *Proceedings of the 2005 IEEE/RSJ International Conference on Intelligent Robots and Systems*, Edmonton, Alberta, Canada. Piscataway, NJ: IEEE. See http://www.urbiforge.com

Barto, A. G., & Şimşek, Ö. (2005). Intrinsic motivation for reinforcement learning systems. In *Proceedings of the 13th Yale Workshop on Adaptive and Learning Systems*, New Haven, CT, (pp. 113–118).

Blank, D., Kumar, D., Meeden, L., & Marshall, J. (2005). Bringing up robot: Fundamental mechanisms for creating a self-motivated, self-organizing architecture. *Cybernetics and Systems*, *36*, 125–150.

Bonarini, A., Lazaric, A., Restelli, M., & Vitali, P. (2006). Self-development framework for reinforcement learning agents. In *Proceedings of the 5th International Conference on Development and Learning (ICDL 2006)*, Bloomington, IN. Piscataway, NJ: IEEE.

Braitenberg, V. (1984). *Vehicles: experiments in synthetic psychology*. Cambridge, MA: MIT Press.

Brooks, A. G., Gray, J., Hoffman, G., Lockerd, A., Lee, H., & Breazeal, C. (2004). Robot's play: interactive games with sociable machines. *ACM Computers in Entertainment*, *2*(3), 10.

Bruner, J. S., & Sherwood, V. (1975). Peekaboo and the learning of rule structures. In J. S. Bruner, A. Jolly, & K. Sylva (Eds.), *Play: Its role in development and evolution*, (pp. 277–285). New York: Penguin.

Clancey, W. J. (1991). Book review: 'The Invention of Memory: A new view of the brain' by Israel Rosenfield. *Artificial Intelligence*, *50*, 241–284.

Clancey, W. J. (1997). *Situated cognition: On human knowledge and computer representations*. Learning in doing: Cognitive and computational perspectives. Cambridge: Cambridge University Press.

Crutchfield, J. P. (1990). Information and its metric. In L. Lam & H. C. Morris (Eds.), *Nonlinear structures in physical systems: Pattern formation, chaos and waves*, (pp. 119–130). New York: Springer.

Dautenhahn, K. (1996). Embodied cognition in animals and artifacts. In *Proceedings of AAAI FS Embodied Cognition and Action*, (pp. 27–32). Technical report FS-96-02. Menlo Park, CA: AAAI Press.

Dautenhahn, K., Bond, A. H., Cañamero, L., & Edmonds, B. (2002). (Eds.). *Socially intelligent agents: Creating rela-*

tionships with computers and robots, Vol. 3 of *Multi-agent systems, artificial societies, and simulated organizations*. Berlin: Springer.

Dautenhahn, K., & Christaller, T. (1996). Remembering, rehearsal and empathy – towards a social and embodied cognitive psychology for artifacts. In S. Ò Nualláin, P. McKevitt, & E. Mac Aogáin (Eds.), *Two sciences of the mind: Readings in cognitive science and consciousness*, (pp. 257–282). Philadelphia, PA: John Benjamins North America.

Glenberg, A. M. (1997). What is memory for? *Behavioral and Brain Sciences*, *20*, 1–55.

Gurney, K. (1997). *An introduction to neural networks*. Boca Raton, FL: CRC Press.

Jaeger, H., & Haas, H. (2004). Harnessing nonlinearity: Predicting chaotic systems and saving energy in wireless communications. *Science*, *304*, 78–80.

Kaplan, F., & Hafner, V. V. (2005). Mapping the space of skills: An approach for comparing embodied sensorimotor organizations. In *Proceedings of the 4th IEEE International Conference on Development and Learning (ICDL 05)*, Osaka, Japan, (pp. 129–134). Piscataway, NJ: IEEE.

Kaplan, F., & Oudeyer, P.-Y. (2007). The progress-drive hypothesis: An interpretation of early imitation. In K. Dautenhahn & C. Nehaniv (Eds.), *Models and mechanisms of imitation and social learning: Behavioural, social and communication dimensions*. Cambridge: Cambridge University Press (pp. 361–377).

Kelso, S. (1995). *Dynamic patterns, the self-organization of brain and behavior*. Cambridge, MA: MIT Press.

Kozima, H., Nakagawa, C., & Yano, H. (2005). Using robots for the study of human social development. In *Proceedings of the AAAI Spring Symposium on Developmental Robotics*, Palo Alto, CA, (pp. 111–114). Menlo Park, CA: AAAI Press.

Lungarella, M., Metta, G., Pfeifer, R., & Sandini, G. (2003). Developmental robotics: A survey. *Connection Science*, *15*, 151–190.

Lungarella, M., & Sporns, O. (2005). Information self-structuring: Key principles for learning and development. In *Proceedings of the 4th International Conference on Development and Learning*, Osaka, Japan, (pp. 25–30). Piscataway, NJ: IEEE.

Matarić, M. J. (1992). Integration of representation into goal-driven behaviour-based robots. *IEEE Transactions on Robotics and Automation*, *8*, 304–312.

Meltzoff, A., & Moore, M. (1997). Explaining facial imitation: A theoretical model. *Early Development and Parenting*, *6*, 179–192.

Michaud, F., & Matarić, M. J. (1998). Learning from history for behavior-based mobile robots in non-stationary conditions. *Machine Learning*, *31*(1–3), 141–167.

Mirza, N. A., Nehaniv, C. L., Dautenhahn, K., & te Boekhorst, R. (2005a). Using temporal information distance to locate sensorimotor experience in a metric space. In *Proceedings of the 2005 IEEE Congress on Evolutionary Computation (CEC2005)*, Edinburgh, Scotland, UK (Vol. 1, pp. 150–157). Piscataway, NJ: IEEE.

Mirza, N. A., Nehaniv, C. L., Dautenhahn, K., & te Boekhorst, R. (2005b). Using sensory-motor phase-plots to characterize robot–environment interactions. In *Proceedings of the 6th IEEE International Symposium on Computational Intelligence in Robotics and Automation (CIRA2005)*, Espoo, Finland, (pp. 581–586). Piscataway, NJ: IEEE.

Montague, D. P. F., & Walker-Andrews, A. S. (2001). Peekaboo: A new look at infant's perception of emotion expression. *Developmental Psychology*, *37*, 826–838.

Nehaniv, C. L. (1999). Narrative for artifacts: Transcending context and self. In P. Sengers & M. Mateas (Eds.), *Narrative intelligence: Papers from the 1999 AAAI Fall Symposium*, North Falmouth, MA, (pp. 101–104). Menlo Park, CA: AAAI Press.

Nehaniv, C. L. (2005). Sensorimotor experience and its metrics. In *Proceedings of the 2005 IEEE Congress on Evolutionary Computation*, Edinburgh, Scotland, UK, (Vol. 1, pp. 142–149). Piscataway, NJ: IEEE.

Nehaniv, C. L., Polani, D., Dautenhahn, K., te Boekhorst, R., & Cañamero, L. (2002). Meaningful information, sensor evolution, and the temporal horizon of embodied organisms. In *Artificial life VIII*, (pp. 345–349). Cambridge, MA: MIT Press.

Nuxoll, A., & Laird, J. E. (2004). A cognitive model of episodic memory integrated with a general cognitive architecture. In M. Lovett, C. Schunn, C. Lebiere, & P. Munro (Eds.), *Proceedings of the 6th International Conference on Cognitive Modeling*, Pittsburgh, PA, (pp. 220–225). Mahwah, NJ: Lawrence Erlbaum Associates.

Olsson, L., Nehaniv, C. L., & Polani, D. (2004). Sensory channel grouping and structure from uninterpreted sensor data. In *2004 NASA/DoD Conference on Evolvable Hardware*, Seattle, WA, (pp. 153–160). Los Alamitos, CA: IEEE Computer Society Press.

Olsson, L., Nehaniv, C. L., & Polani, D. (2005). Discovering motion flow by temporal-informational correlations in sensors. In L. Berthouze, H. Kozima, C. G. Prince, G. Sandini, G. Stojanov, G. Metta, & C. Balkenius (Eds.), *Proceedings of the 5th International Workshop on Epigenetic Robotics: Modeling Cognitive Development in Robotic Systems* (pp. 117–120). Lund University Cognitive Studies, 123. Lund, Sweden: LUCS.

Olsson, L., Nehaniv, C. L., & Polani, D. (2006a). From unknown sensors and actuators to actions grounded in sensorimotor perceptions. *Connection Science*, *18*, 121–144.

Olsson, L., Nehaniv, C. L., & Polani, D. (2006b). Measuring informational distances between sensors and sensor integration. In L. M. Rocha, L. S. Yaeger, M. A. Bedau, D.

Floreano, R. L. Goldstone, & A. Vespignani (Eds.), *Artificial life X*. Cambridge, MA: MIT Press.
OpenCV. (2000). Open computer vision library (gpl licence). See http://sourceforge.net/projects/opencvlibrary/
Oudeyer, P.-Y., Kaplan, F., Hafner, V. V., & Whyte, A. (2005). The playground experiment: Task-independent development of a curious robot. In D. Bank & L. Meeden (Eds.), *Proceedings of the AAAI Spring Symposium on Developmental Robotics*, Palo Alto, CA, (pp. 42–47). Menlo Park, CA: AAAI Press.
Pea, R. D. (2004). The social and technological dimensions of scaffolding and related theoretical concepts for learning, education and human activity. *Journal of the Learning Sciences*, *13*, 423–451.
Pfeifer, R., & Scheier, C. (1999). *Understanding intelligence*. Cambridge, MA: MIT Press.
Pierce, D., & Kuipers, B. (1997). Map learning with uninterpreted sensors and effectors. *Artificial Intelligence*, *92*, 169–229.
Quick, T., Dautenhahn, K., Nehaniv, C. L., & Roberts, G. (1999). The essence of embodiment: A framework for understanding and exploiting structural coupling between system and environment. In *Proceedings of the 3rd International Conference on Computing Anticipatory Systems (CASYS'99)*, Liège, Belgium. AIP Conference Proceedings.
Rochat, P., Querido, J. G., and Striano, T. (1999). Emerging sensitivity to the timing and structure of protoconversation in early infancy. *Developmental Psychology*, *35*, 950–957.
Rolls, E. T. (1999). *The brain and emotion*. Oxford: Oxford University Press.
Rosenfield, I. (1988). *The invention of memory: A new view of the brain*. New York: Basic Books.
Rylatt, R. M., & Czarnecki, C. A. (2000). Embedding connectionist autonomous agents in time: The 'road sign problem'. *Neural Processing Letters*, *12*, 145–158.
Scheult, M., & Logan, B. 2001. Affective vs deliberative agent control. In *Proceedings of AISB 2001*, University of York, UK, (pp. 1–10). York, UK: Society for the Study of Artificial Intelligence and Simulation of Behaviour.
Shannon, C. E. (1948). A mathematical theory of communication. *Bell Systems Technical Journal*, *27*, 379–423 and 623–656.
Sutton, R. S., & Barto, A. G. (1998). *Reinforcement learning: An introduction*. Cambridge, MA: MIT Press.
Tarapore, G., Lungarella, M., & Gómez, G. (2004). Fingerprinting agent–environment interaction via information theory. In F. Groen, N. Amato, A. Bonarini, E. Yoshida, & B. Kröse (Eds.), *Proceedings of the 8th International Conference on Intelligent Autonomous Systems*, Amsterdam, the Netherlands, (pp. 512–520). IOS Press.
te Boekhorst, I. J. A., Lungarella, M., & Pfeifer, R. (2003). Dimensionality reduction through sensori-motor coordination. In O. Kaynak, E. Alpaydin, E. Öja, & L. Xu (Eds.), *Proceedings of the Joint International Conference on Artificial Neural Networks and Neural Information Processing*, Istanbul, Turkey, (pp. 496–503). *Lecture Notes in Computer Science*, Vol. 2114. Berlin: Springer.
Thelen, E., & Smith, L. B. (1994). *A dynamic systems approach to the development of cognition and action*. Cambridge, MA: MIT Press.
Tulving, E. (1983). *Elements of episodic memory*. Oxford: Clarendon Press.
van Duijn, M., Keijzer, F., & Franken, D. (2006). Principles of minimal cognition: Casting cognition as sensorimotor coordination. *Adaptive Behavior*, *14*, 157–170.
Veatch, T. C. (1998). A theory of humour. *International Journal of Humour Research*, *11*, 161–175.
von Hofsten, C. (1993). Prospective control: A basic aspect of action development. *Human Development*, *36*, 253–270.
Vygotsky, L. S. (1978). *Mind and society: The development of higher mental processes*. Cambridge, MA: Harvard University Press.

About the Authors

Naeem Assif Mirza received a B.Sc. with honors in microelectronic systems engineering from the University of Manchester Institute of Science and Technology, Manchester, UK, in 1992. After a career in telecommunications and investment banking, he received an M.Sc. with distinction in artificial life from the University of Hertfordshire, Hatfield, UK, in 2004. He is now a Ph.D. candidate and research assistant at the University of Hertfordshire, working with the RobotCUB European project, researching interaction histories for ontogeny in robots. He is a member of the euCognition European Network for the Advancement of Artificial Cognitive Systems.

Chrystopher L. Nehaniv received a B.Sc. with honors in mathematics (University of Michigan, Ann Arbor, 1987), a Ph.D. in mathematics (University of California, Berkeley, 1992) for work in the algebraic theory of semigroups, groups, and automata, and was a lecturer and research fellow at Berkeley until 1993. He was a Professor of the University of Aizu, Japan, and Director of the Software Engineering Lab until 1998. He is a visiting Professor at the Institute for Mathematics and Informatics, University of Debrecen, Hungary, and has taught at Ibaraki National University in Japan. In 1999 he became a reader in computer science and since 2001 a research professor of mathematical and evolutionary computer science, with the Interactive Systems Engineering Research Group, Adaptive Systems Research Group, and the Algorithms Research Group at the University of Hertfordshire, UK. He is Director of the UK EPSRC Network on Evolvability in Biological and Software Systems, Associate Editor of the Elsevier journal *BioSystems: Journal of Biological and Information Processing Sciences*, Associate Editor of *Interaction Studies: Social Behaviour and Communication in Biological and Artificial Systems* and a member of the Santa Fe Institute Evolvability Working Group. *E-mail*: c.l.nehaniv@herts.ac.uk

Kerstin Dautenhahn is a research professor in the School of Computer Science (Faculty of Engineering and Information Sciences) at University of Hertfordshire, UK, where she is coordinator of the Adaptive Systems Research Group. Her main areas of research are human–robot interaction, social robotics, socially intelligent agents, and artificial life. She is a former member of the Department of Biological Cybernetics at the University of Bielefeld, Germany (1990–1993), the AI-Lab at GMD, Sankt Augustin, Germany (1993–1996), and the VUB Artificial Intelligence Laboratory, Brussels, Belgium (until end of 1996). From January 1997 to April 2000 she was a lecturer at the Department of Cybernetics, University of Reading, UK. *E-mail*: k.dautenhahn@herts.ac.uk

Rene te Boekhorst is a senior lecturer at the School of Computer Science, University of Hertfordshire, UK. He was trained as a biologist at the University of Utrecht (the Netherlands) and has a special interest in applied statistics, complex systems theory and computational biology. For his Ph.D. thesis he applied individual-oriented models to explain field data on the behavioral ecology and social ethology of great apes in terms of processes of self-organization. As a research assistant at the University of Zurich, he used ethological methods to study the collective behavior of robots. Currently, he is active in the area of robotics (mainly concentrating on swarm intelligence and human–robot interactions), dynamic systems modeling of motivation, emotion and behavior and statistical applications in bioinformatics and genomics. *E-mail*: r.teboekhorst@herts.ac.uk

Interaction History Architecture and Scalability**

Naeem Assif Mirza, Chrystopher L. Nehaniv

September 26, 2007

Abstract

This technical article discusses issues of scalability in the interaction history architecture.

1 Scalability

1.1 Introduction

In this section we discuss scalability of the interaction history architecture in terms of computation time and memory requirement as the history of the robot's interaction with its environment grows. The specific areas which can result in scalability issues are 1) the time to place a new experience in a metric space (for the purpose of returning a list of nearest neighbours) and 2) the storage of the metric space of experiences. The goal is that computation time should allow "real-time" operation (*i.e.* the computational complexity should be constant) and that there should be mechanisms to keep within a fixed memory limit.

We look at methods that reduce the processing requirement for placing new experiences, as well as forgetting of experiences and the merging of experiences.

1.2 Computational Complexity

Placing an experience in a metric space requires that the distance from that experience to every other one in the space be known. By far the most computationally expensive task in this process is the calculation of Experience Distance between any two experiences. Without any modifications to the basic architecture, the time to place an experience increases linearly with the number of experiences already in the space. All other processing requirements are constant.

Given that new experiences arrive regularly, it is inevitable that as the number of experiences in the space grows it will not be possible to place an experience in the metric space before another one is available for processing. Further, given that the time to make a single comparison is constant, the only way to reduce computation time is to reduce the number of comparisons. This can be done either by reducing the number of items to compare (see Section 1.3) or by not explicitly computing all distances (see Section 1.2.1).

1.2.1 Reducing Comparisons

One way to reduce required comparisons would be to use the distances between experiences in the space to infer distances to any new experience. However, as the metric space of experiences is a non-euclidean space, then this becomes more difficult.

We note, however, that knowing the distance to all other experiences is not necessary for the correct operation of the interaction history architecture. It is only necessary to know the *nearest* neighbours; *i.e.* the nearest N experiences, or all experiences within a "ball" of radius r.

**The work described in this paper 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. We are grateful to Kerstin Dautenhahn and René te Boekhorst for discussions on this work.

1.2.2 Finding Nearest Neighbours

Say we are interested in finding all nearest neighbours of an experience E^{new} within a "ball" of radius r, we can use the triangle inequality to reduce the number of distances that need to be measured. Specifically:

Theorem 1.1 *Given an experience E^k that is distance $d(E^{new}, E^k) \leq r$ from E^{new}, then any neighbours of E^k that are further away than $2r$ are not within distance r of E^{new}.*

Proof 1.1 *Consider 2 experiences Y, Z near X; We define near to mean within distance r, thus: $d(X, Y) \leq r$ and $d(X, Z) \leq r$. Then by the triangle inequality $d(Y, Z) \leq d(X, Y) + d(X, Z)$, we get $d(Y, Z) \leq 2r$. Therefore, if any 2 experiences are further apart than $2r$, then they cannot both be within radius r of any particular experience.*

This fact can be used discard experiences from consideration when finding nearest neighbours within a specified radius. Of course, this requires first finding an experience with radius r of the new experience. One approach to this problem is to simply randomly sample the experience space until one is found. Other strategies exist; for example, using the continuous nature of the environment to start the search for near experiences (in terms of information distance) with those experiences near in terms of time.

Algorithm 1 Algorithm `B2R_NN`: Populate Metric Space Distances for Nearest Neighbours

Require: r = radius for near experiences
Require: E^{new} = new experience
Require: $newDistances$ = empty list

```
toTestList ⇐ all experiences in metric space
while toTestList is not empty do
   remove a random experience from toTestList, assign to E^K
   calculate d(E^new, E^k), add to newDistances
   if D ≤ R then
      remove all experiences further than 2r from E^k in toTestList
   end if
end while
```

An algorithm to find the nearest neighbours of a new experience from a metric space of experiences is given in Algorithm 1 which guarantees that all experiences within r of the new experience E^k will be in the list $newDistances$. There will also be have been other distances which will have been checked in that list as a consequence of the random sampling.

An important issue is that any strategy that does not fully populate all distances in a metric space is potentially degenerative. That is, when another experience arrives, it may not be possible to make the same guarantees as the existing metric space is not fully populated. In practical use however, the algorithm given should still find all neighbours as it excludes only experiences which clearly do not fall within radius r. The result instead is that potentially more comparisons will have to be made. This however in turn results in a better populated space.

A question remains: by how much this might reduce the space of experience to be searched? The answer is largely dependant on r (as shown in the tests below, see Section 1.2.4) and on the nature of the space. At one extreme, if experiences are clustered tightly together with no experience further than $2r$ from any other, then all experiences must be searched. Due to the nature of the algorithm, the computation time would actually be greater than if all experiences were checked in turn. At the other extreme, if the radius was smaller than any distance between two experiences, then once again all experiences would have to be checked because no near neighbour would be found.

Happily, the situation is likely to be somewhere between the two. If the experiences are clustered around many centres further apart than $2r$, or are evenly spaced with the minimum distance much less than r but the maximum distance much greater than $2r$, then it is likely that a near experience will be found fairly quickly and consequently, many experiences will be discarded, reducing the computation time significantly.

1.2.3 Finding a suitable radius r

With the strategy given above, an important question is: what value should r take? This clearly depends on the nature of the space and how many nearest neighbours are needed to make a choice of next action within the interaction history architecture action selection. Thus, r is likely to change as the robot interacts in the environment and so should be adaptive.

A strategy to adapt r suitably to the current metric space is to instead take N, the number of nearest neighbours required, as a reference point. Starting with r at an initial value, for every new experience, find all neighbours within radius r. If this number is greater than N, adjust r downwards and *visa-versa*.

1.2.4 Test of `B2R_NN` algorithm in artificial and real metric spaces

To quantify the computational saving that can be achieved by the `B2R_NN` algorithm, we conducted two tests. Firstly, an artificial euclidean metric space with evenly spaced random points was used to investigate the relationship between the density of the points in the space, N, and the radius, r. Secondly, a real metric space of experience taken from an Aibo interacting with a human partner was used to investigate the effect of varying the radius, r.

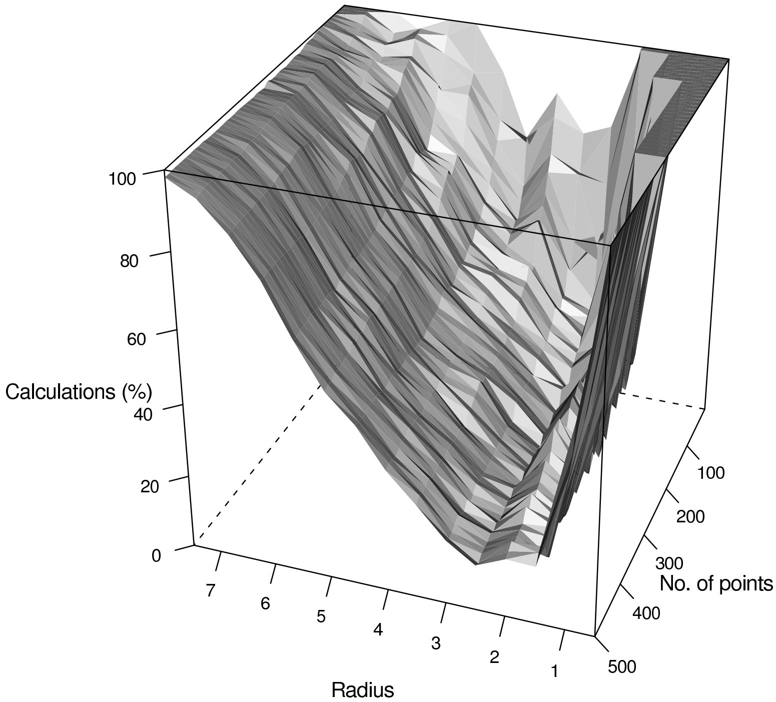

Figure 1: Graphic showing relationship between r and N in determining the number of calculations made by the `B2R_NN` algorithm in an artificial metric space.

In figure 2 we show the results from the artificial space. The metric space was 3-dimensional euclidean and contained randomly placed points. The maximum possible distance in the space was 17.32 *(no units)*, with an observed average distance between any 2 points of approx. 7.2 and a minimum distance to any neighbour between 2.3 and 0.7 depending on the number of points.

The results show that when the radius r is relatively small (in this case $r \leq 1.0$) then there is no or very little reduction in the number of calculations required to find the neighbours in a ball of radius r. As the radius increases, less than 20% of the calculations are needed, However, this saving is lessened as the radius grows until it eventually comes back down to 0. While these observations are true to some extent for any number of points, the certainty of gaining such a speed-up is increased with the density of points in the space.

Figure 3 shows the results when the algorithm was testing in a metric space that resulted from an Aibo interacting with a human partner. The Aibo variously looked at the partner's face, hid its face with it's forearm (peekaboo) and looked at the pink ball. The space had a total of 372 experiences in it. The distances for the 373^{rd} experience were pre-calculated for the purposes of the test, and used as a look-up table in the tests of the `B2R_NN` algorithm.

A similar shaped curve is again observed indicating that, with a good choice of r, significant saving in number of calculations can be achieved.

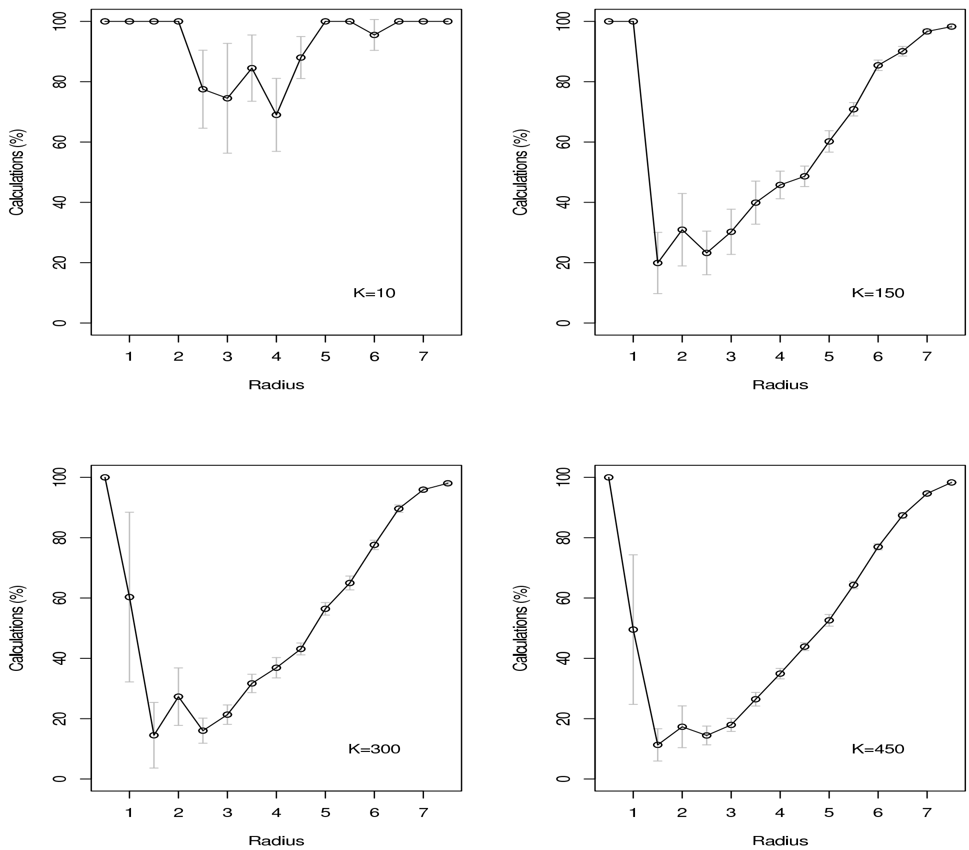

Figure 2: Graphs showing relationship between r and N in determining the number of calculations made by the B2R_NN algorithm in an artificial metric space for selected N. Each point is the mean of 20 runs, error bars show 1 *Std. Dev.*

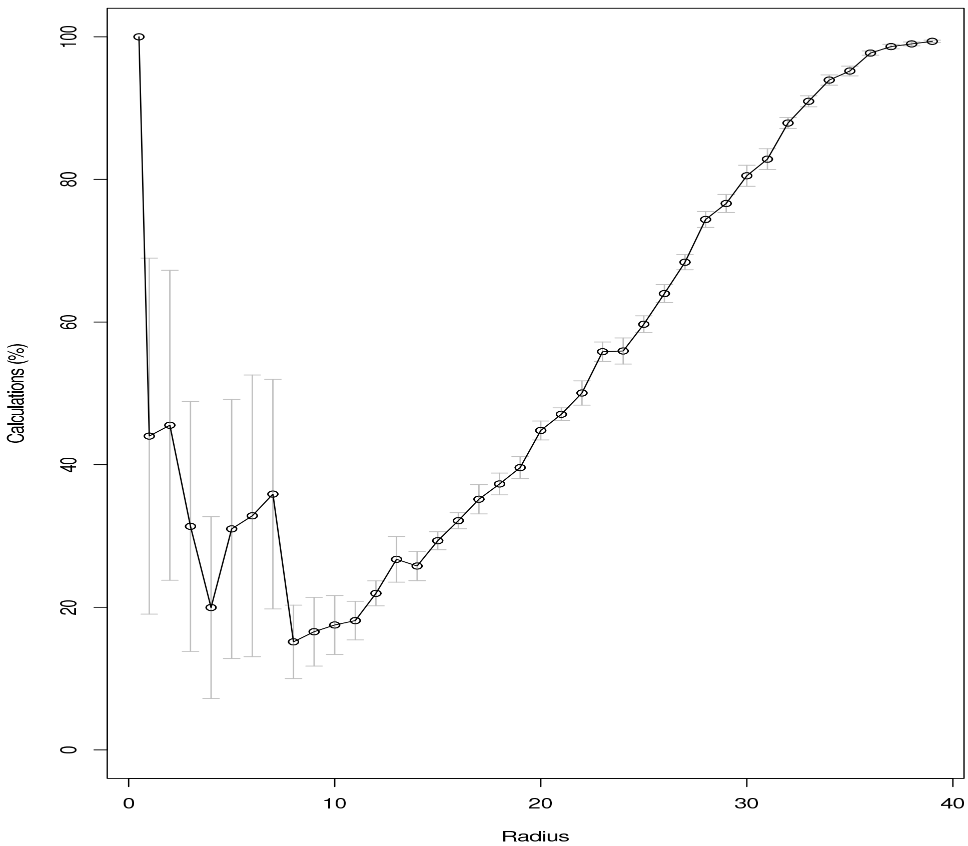

Figure 3: Graph showing effect of r in determining the number of calculations made by the `B2R_NN` algorithm in a real metric space taken from an Aibo interacting with a human partner. Each point is the mean of 20 runs, error bars show 1 *Std. Dev.*

1.2.5 Reducing number of points

Another strategy for reducing the number of computations of distance between experiences is to reduce the number of experiences in the space in the first place. This is the subject of the next section, and the discussion continues there.

1.3 Storage Requirements

The memory storage required to maintain a experience space consists of: the storage of the experience[1] itself, plus that of the metric space itself (*i.e.* distances), plus constant factors. The storage of fixed length experiences grows linearly with the number of experiences. The non-constant storage of the metric space itself increases faster. At any time it is proportional to P_2^N, the number of permutations of 2 items from N items, where N is the number of experiences. In terms of complexity this is order $O\left(n \log n\right)$.

Thus, it is not possible to store all experiences and all distances indefinitely for a metric space that is growing. At some point it will exceed the storage available. Also, many calculations in the space are dependant on the number of experiences and so computational complexity is also affected. We therefore examine two strategies to reduce the number of experiences within a metric space as it is growing: *merging* and *forgetting*.

1.3.1 Merging Experiences in a Growing Metric Space

This strategy is based on the idea that if two experiences are very similar, then they could potentially be treated as the same experience for the purposes of comparison with other experiences. Intuitively, this is what happens with us as we experience the world. As we engage in an activity that we do many times, such as drinking a mug of tea at our desk, we do not notice that it is similar to any one particular time we engaged in that activity in the past, only that it is similar to a generalized activity: *i.e.* past experiences have been merged into a single experience (for the purpose of comparison at least).

The general strategy is to replace two experiences in the space by a single experience that has features taken from both. Individual strategies are distinguished by how the two experiences are chosen, *e.g.* by using a threshold T^{merge}, and by what features of the experiences are retained or discarded.

1.3.2 On Calculating an Intermediate Experience

Merging two experiences E^a and E^b, and replacing them with one that is some-way between the two, can be considered as a problem of finding an *intermediate* experience E^{ab}, such that: $d(E^{ab}, E^a) \leq d(E^a, E^b) \wedge d(E^{ab}, E^b) \leq d(E^a, E^b)$.

Ideally the intermediate experience would be half-way between the two, *i.e.* $d(E^{ab}, E^a) = d(E^{ab}, E^b)$. This calculation is not mathematically straightforward due to the non-euclidean nature of the space and may take quite long to compute. One possibility is to find a combination of binned sensor readings that is approximately half the hamming distance between the two sets of values, however ee do not explore this possibility further here.

Alternatively, one or other of the distances can be zero, which amounts to keeping one of the experiences and removing the other. This becomes less of a problem as $d(E^a, E^b)$ approaches zero.

1.3.3 Merging by deletion

Due to the difficulty of mathematically merging two experiences an alternative strategy is to remove one of the experiences entirely. This may not be satisfactory as that experience probably had important information that may be useful. The fact of its existence is one such. *i.e.* the fact that it occurred and was similar to other experiences gives a sense of *familiarity* and may be important in choosing a list of N nearest neighbours. Another important piece of information is the subsequent action that was taken after that particular experience, which may or may not have been different from the other experience. Finally, the distance information may also be important.

A modified strategy would be to remove one of experiences from the space, but retain other information such as number of merged experiences and subsequent actions with the remaining experience. This is in fact the preferred strategy in our architecture. See Algorithm 2 for details.

[1] In storing an experience, all that is required are the binned values of the sensors, not the actual values of the sensors. In addition meta-information will be stored with the experience *e.g.* next action, quality, weight *etc.*

1.3.4 Choice of Experiences to Merge

An obvious choice for merging criteria is to merge any two experiences closer than a threshold T^{merge} (see Algorithm 2. T^{merge} could be fixed but that raises the problem of finding a suitable value. Alternatively it could be an adaptive threshold responding to some other criteria such as maximum number of experiences in the space. For the special case $T^{merge} = 0$ no information is lost in the merge of the sensorimotor experiences themselves, however, they can still be different actual experiences, with different meta-variables (e.g. subsequent actions, quality *etc.*) attached.

An alternative way of choosing experiences to merge would be to compare other features of the experiences. Candidates are *next action* and assigned *quality*. Of course, these other features could be combined with threshold to refine the choice.

Algorithm 2 Algorithm `IHA_MERGET`: Choose and Merge 2 experiences using a threshold

for E^i **in** all experiences **do**
  for E^j **in** neighbours of E^i **do**
    if $d(E^i, E^j) \leq T^{merge}$ **then**
      $actions(E^i) = actions(E^i) + actions(E^j)$
      $quality(E^i) = \left(quality(E^i) + quality(E^j)\right)/2$
      $weight(E^i) = weight(E^i) + weight(E^j)$
      delete all distances to and from E^j in the metric space
      delete E^j
    end if
  end for
end for

1.3.5 Retaining Distances

An alternative to the complete removal of an experience from the metric space, is to delete only the sensorimotor experience data and retain only the existing distance information. This will result in a reduction of memory requirement while retaining important structural information about the metric space. The space would then contain *parent* experiences about which everthing is known, and *child* experiences having only distance and meta information. Any new experiences would only be able to be directly compared to parents and distances to child experiences only inferred.

This strategy has the advantage that only sensorimotor information is lost, and that a natural hierarchy within the metric space can easily be built. The disadvantage is in that the distances from child experiences cannot be known, and that the complexity and storage requirements is not reduced significantly.

1.4 Forgetting

In terms of a metric space of experiences, *forgetting* corresponds to removing individual experiences from the space, including all meta-information and distances to other experiences. This is a useful way to reduce both computational complexity in maintaining the space as well as reducing storage requirement.

The question is; how should experiences be chosen for removal? Of course, it could be random, however it makes more sense to base this on some quality of the experience itself. For instance, time. *i.e.* how often the experience has been “accessed” or “used” or how long ago it was last “accessed”. This would correspond to natural, intuitive ideas of forgetting. An alternative measure could be the quality of the experience in terms of reward signals. In this scheme, experiences that were neither “very good” nor “very bad” might be candidates for forgetting. In terms of the metric space of experiences, another measure might be how isolated and experiences is from others.

Interaction History Architecture : A Simple Test-bed**

Naeem Assif Mirza, Chrystopher L. Nehaniv

September 26, 2007

Abstract

This technical article discusses experiments with the interaction history architecture on a toy learning problem.

1 Introduction

In the classic T-Maze task, an agent (*e.g.* rat or wheeled robot) is required to navigate a simple mazes with a reward at the end of one arm of the T. Also known as a *delayed response task*, this is a popular test-bed for reinforcement learning as the reward is given at sometime after the decision to turn left or right at the junction is taken. The Road-Sign problem is an extension of a simple T-Maze learning environment where an indication of the reward position is given by an earlier disconnected event. Thus the agent can make use of its experience in making the decision to turn left or right. This problem provides a benchmark test-bed for autonomous agents with some kind of short-term memory.

Up till now the interaction history architecture has been applied to complex interaction tasks involving a real robot with multiple degrees of freedom and colour vision (Sony's *Aibo*). The two main test environments were the prediction of the path of a ball using previous experience and developing the capability to play a modified version of the peekaboo interaction game. A simpler benchmark problem however may reveal as yet unexamined capabilities and deficiencies in the architecture and allows the opportunity to extend and develop the architecture to include, for example, adaptive horizon lengths. We also use this simple environment to test the information distance measure against other measures such as hamming distance.

1.1 The Road-Sign Problem

The agent is in a T-Maze at the bottom end of the T. The basic task is to travel to the junction and turn either right or left. A reward is placed at one arm of the T. Which particular arm is a variable of the experiment.

While travelling to the junction, the agent encounters a signal in the form of a light on either the left-hand side or the right-hand side. In the simplest version of the task this faithfully indicates the position of the reward in the T. Of course more complex relationships between signal and reward can be devised.

In some versions of the experiment, a negative reward may be given if the agent travels to the end of the "wrong" arm of the T.

1.2 Aims

The aim of this study is to investigate how the interaction history operates in a simple benchmark test and in particular to address the following questions:

Learning : How well does the system perform in this simple task, *i.e.* is it able to associate the signal and reward over a series of runs through the maze?

Alternative Distance Measures : How does the information distance metric compare to alternative measures of distance such as the hamming distance and the pearson's squared metric?

Multiple Horizon Lengths : Using multiple metric spaces maintained simultaneously with different horizon lengths is the system able to choose actions from experiences with appropriate horizon length?

**The work described in this paper 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. We are grateful to Kerstin Dautenhahn and René te Boekhorst for discussions on this work.

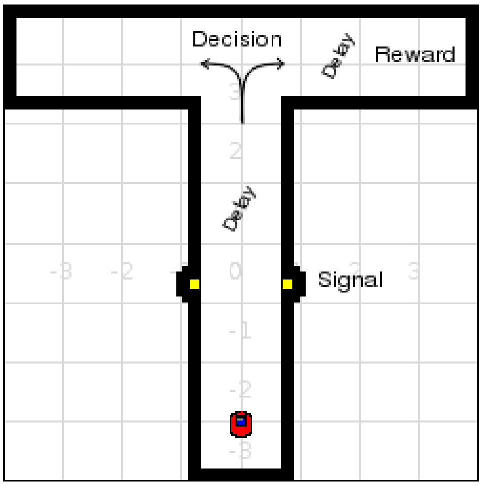

Figure 1: The T-Maze task

2 Related Work

Much of the recent literature on solutions to this problem for autonomous agents are either neural-network based or evolutionary algorithm based.

Rylatt and Czarnecki (2000) describe a Elman-style recurrent neural network solution using a type of learning called CRBP *Complementary Reinforcement Backpropagation Learning.* Although in that original paper they do not tackle the whole problem. Thieme and Ziemke (2002) go further, testing four different neural network architectureson the road-sign problem with the highest reliability achieved by *Extended Sequential Cascaded Networks*, a higher-order recurrent neural network architecture. They showed that a short-term memory can be realized for delayed response tasks through synaptic plasticity and dynamic modulation of sensorimotor mapping.

Interestingly, Thieme and Ziemke (2002) also found that a simple feed-forward neural network could also reactively solve the road-sign problem. This is achieved by moving towards the light and then simply following the wall till the goal is reached. The memory of state is in the agent-environment interaction. In a similar vein, Bovet and Pfeifer (2005) explore the possibility that a memory-less agent could solve the road-sign problem. In their case, appearance of a short-term memory is achieved through a combination of some unchanging aspect of the environment (a coloured wall) and plasticity of synaptic weights between reward and the visual modality. In effect, the visual system has been altered by the interaction with the light and the subsequent presentation of the environmental stimulus induces the appropriate motor response.

Kim (2004) takes an alternative approach of evolving a controller based on Finate State Machines to analyze the role of internal memory. They looked the size of the internal memory and states required to learn various forms of the problem involving one, two or more lights. They also studied the effect of noise on their model. They find that purely reactive controllers cannot solve the problem and multi-states were required. The simplest arrangement of a single light requiring two-states, with more states required as the number of lights increased.

The representation of state is approached in many different ways in the above. Linåker and Jacobsson (2001) work at a high level, extracting significant events and clustering them to reduce the number of states down to a handful. They use a vector quantization network to extract model vectors representing event classes. These provide inputs to a simple recurrent neural network which learns the associations between events and behaviours.

3 Implementation

Player/Stage was chosen to simulate a robot and the maze itself. The software would be written entirely in C++ under the YARP framework. The simulation uses a pioneer robot model with a SICK laser scanner for localization, and a CMU camera with colour blob detection in the place of vision.

The robot collects sensorimotor data continuously creating experiences and placing them in a metric space. For this implementation the system allows for spaces of multiple horizon lengths to be built on-line simultaneously. The experimental runs consist of multiple iterations of a maze with different positions of lights and reward with the robot being placed back at the start with its history intact after it has reached one or other end of the T.

3.1 Reward

The motivational system is a simple reward signal and returns 1 when the robot reaches the end of the correct arm of the T and 0 at all other times. This scheme is used in these experiments. Alternative schemes can have negative rewards for reaching the wrong end of the arm, as well as returning an intermediate value while the robot is traversing the maze.

3.2 Actions

In order to study the effect of the interaction history in detail the robot is contrained to make a single action selection decision (turn left or right) at the junction of the T. In exploratory trials the system was less contrained, but this led to difficulties interpreting the results so the situation was simplified to have a single decsion point that could be compared across trials.

4 Experimental Methodology

During early testing, a methodology similar to that used for the Aibo "peekaboo" experiments was followed. Specifically, that each trial run was started with an empty history of experience, and experience was gathered on-line by using random exploration. However, this results in experimental results that are hard to compare as they have different histories on which to base decisions. As a result an alternative strategy was used whereby a common history was used across repeated runs of a particular trial alowing a fair comparison between trials. The common history was gathered during a single run where the robot was constrained to make the correct decisions.

5 Summary of Results

5.1 Learning

The first result was that it was shown that, given the right conditions, the robot was able to take the correct turn on 19 out of 20 decision points across 10 trials using a common history. The conditions were that a single horizon was chosen which was long enough to cover the period from the light to the end of the arm of the T. Where the horizon was too short, the robot was not able to learn the task. Secondly, the choice of nearest neighbours was reduced to two. More nearest neighbours resulted in, randomly, incorrect decisions being taken as the common history was short and only had two examples of each turn.

The common history contained 219 experiences over 4 iterations of the maze, with the reward alternating between left and right over the 4 scenarios. Figure 2 shows the experience distances from $Exp : 18$, the first decision point (a left turn scenario). It can be seen that the nearest experiences (distance 0.11731) are those around $Exp : 136$ in the 3^{rd} iteration which is also a left turn scenario. The experiences at the decision points in the other scenarios are not as close. Table 1 shows the roulette selection table as created by the action selection process during a subsequent left-turn scenario. It shows that the left turn (action 2) experiences $Exp : 136$,$Exp : 19$ and $Exp : 18$ are the most likely to be selected.

5.2 Alternative Distance Measures

The information distance measure was compared with two other measures of distance, the Hamming metric and the Pearson's Squared Correlation distance. The Hamming distance for numeric quantities is the 1-norm distance (also known as Manhattan distance) and is given by equation 1. This is calculated for the binned sensor values.

$$d_{hamming}(x, y) = \sum |x_i + y_i| \tag{1}$$

The Pearsons correlation r_{xy} is a statistical measure of the correlation between two random variables x and y and can be calculated from samples of x and y using the equation 3. A distance measure based on the cooefficient of determination r^2 can also be calculated (equation 2). This quantity has is interesting for us as it is small for both correlated and anti-correlated variables, as is the information metric.

$$d_{pearsons}(x, y) = 1 - r_{xy}^2 \tag{2}$$

where

$$r_{xy} = \frac{n \sum x_i y_i - \sum x_i \sum y_i}{\sqrt{n \sum x_i^2 - (\sum y_i)^2} \sqrt{n \sum y_i^2 - (\sum y_i)^2}} \tag{3}$$

Figures 4 and 3 show the Hamming and Pearson's distances from experience $Exp : 18$ to all others in the common history as was shown for the information metric in Figure 2. All the measures clearly show most similarity between equivalent experiences (*i.e.* $Exp : 136$ the other turn-left experience). They also show similarity to experiences at the same point in the maze but with the light on the opposite side of the wall. The Pearson's and information metric also show marked similarity between $Exp : 18$ and certain others in the history, showing that they both reveal correlations in the experience beyond the obvious.

5.2.1 Multiple Horizon Lengths

Following the experiments with a fixed single horizon length for experiences, trials were carried out using multiple simultaneous metric spaces of different horizon length experiences. At any action selection point, the system could choose from similar experiences both within a single space as well as from other spaces. It was expected that the choice of experience would reflect the ideal horizon length for the problem at hand. However, instead it was found that the nearest neighbours were consistently of shorter horizon lengths as there is naturally less variation in shorter samples. Thus, when a set of horizons included a horizon length too short to learn the task, the system tended to choose experiences from that metric space and so failed to learn the task. See Figure 5.

In order for this strategy to succeed, it may be necessary to bias the experience choice to favour longer horizons over shorter ones.

Table 1: Example of choice of experiences and asociated actions ordered by weighted distance.

Exp	Hor	Weighted %	Distance	Mass	Value	Action Freq.			
						0	1	2	3
136	64	17.421436%	0.173884	1	1.0	0.0	0.0	1.0	0.0
19	64	12.480776%	0.205438	1	1.0	0.0	0.0	1.0	0.0
18	64	9.317441%	0.237768	1	1.0	0.0	0.0	1.0	0.0
194	64	6.884041%	0.276618	1	1.0	0.0	0.0	0.0	1.0
137	64	5.582502%	0.307176	1	1.0	0.0	0.0	1.0	0.0
120	64	5.166194%	0.319313	1	1.0	0.0	1.0	0.0	0.0
178	64	5.056741%	0.322750	1	1.0	0.0	1.0	0.0	0.0
77	64	4.804414%	0.331117	1	1.0	0.0	0.0	0.0	1.0
78	64	4.593748%	0.338624	1	1.0	0.0	0.0	0.0	1.0
3	64	4.492898%	0.342404	1	1.0	0.0	1.0	0.0	0.0
128	64	3.812116%	0.371722	1	1.0	0.0	1.0	0.0	0.0
62	64	3.770775%	0.373755	1	1.0	0.0	1.0	0.0	0.0
61	64	3.516556%	0.387029	1	1.0	0.0	1.0	0.0	0.0
186	64	3.487553%	0.388635	1	1.0	0.0	1.0	0.0	0.0
121	64	3.224787%	0.404158	1	1.0	0.0	1.0	0.0	0.0
2	64	3.222349%	0.404311	1	1.0	0.0	1.0	0.0	0.0
11	64	3.165672%	0.407914	1	1.0	0.0	1.0	0.0	0.0
251	64	0.000000%	0.374959	1	0.0	1.0	0.0	0.0	0.0
244	64	0.000000%	0.337201	1	0.0	0.0	1.0	0.0	0.0
236	64	0.000000%	0.272928	1	0.0	0.0	1.0	0.0	0.0

Columns: *Exp*: experience number, *Hor*: horizon length, *Weighted %*: chance of selection of experience based on distance, value ad weight, *Distance*: experience distance from current experience, *Mass*: number of merged experiences, *Value*: future expected reward, *Action Freq*: a frequency distribution of next actions from this experience. (Actions are 0=none, 1=Forward, 2=Left, 3=Right)

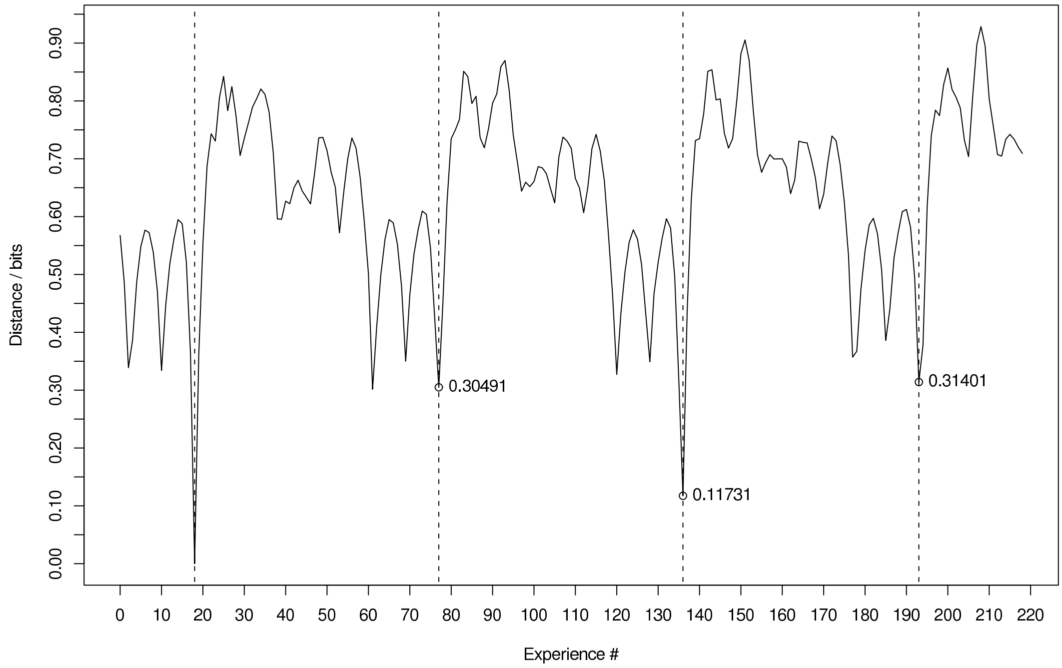

Figure 2: Distances from experience #18 (1st decision point) in common history.

References

S. Bovet and R. Pfeifer. Emergence of delayed reward learning from sensorimotor coordination. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 841–846, 2005.

F. Linåker and H. Jacobsson. Learning delayed response tasks through unsupervised event extraction. *International Journal of Computational Intelligence and Applications*, 1(4):413–426, 2001.

R. M. Rylatt and C.A. Czarnecki. Embedding connectionist autonomous agents in time: The 'road sign problem'. *Neural Processing Letters*, 12(2):145–158, Oct 2000.

Mikael Thieme and Tom Ziemke. The road sign problem revisited: handling delayed response tasks with neural robot controllers. In *ICSAB: Proceedings of the seventh international conference on simulation of adaptive behavior on From animals to animats*, pages 228–229, Cambridge, MA, USA, 2002. MIT Press. ISBN 0-262-58217-1.

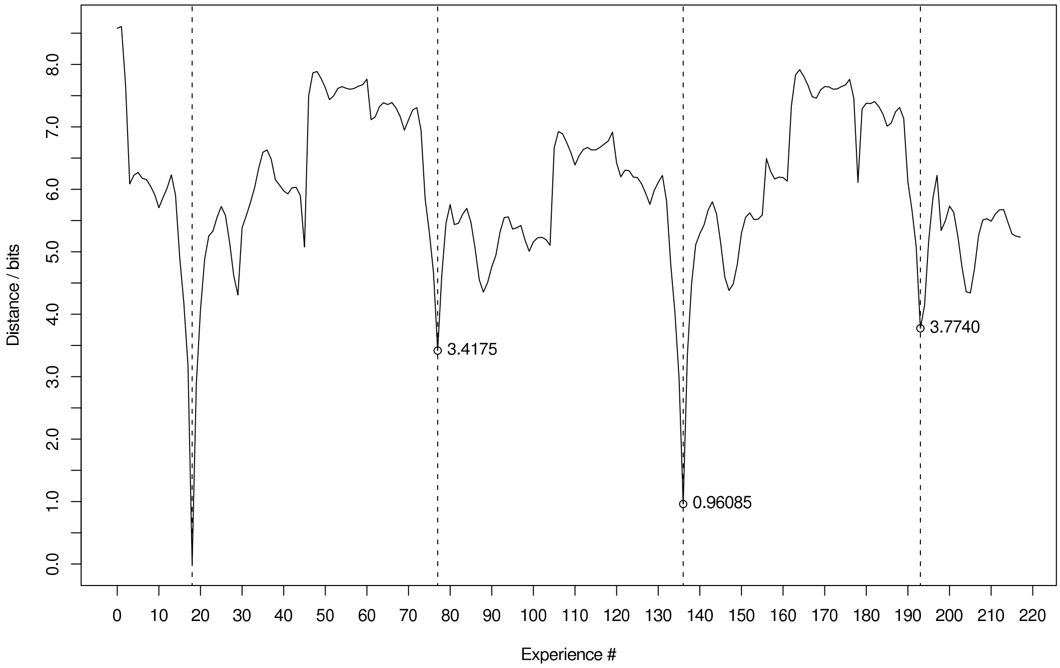

Figure 3: Pearson correlation distances (see Figure 2 for comparison with information distance.

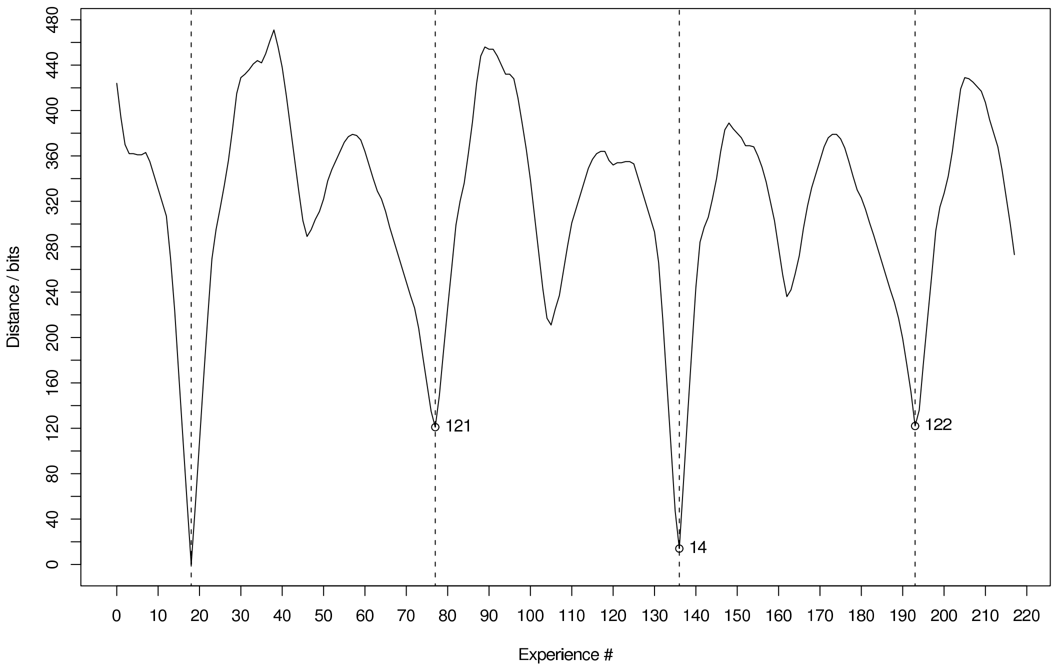

Figure 4: Hamming distances (see Figure 2 for comparison with information distance.

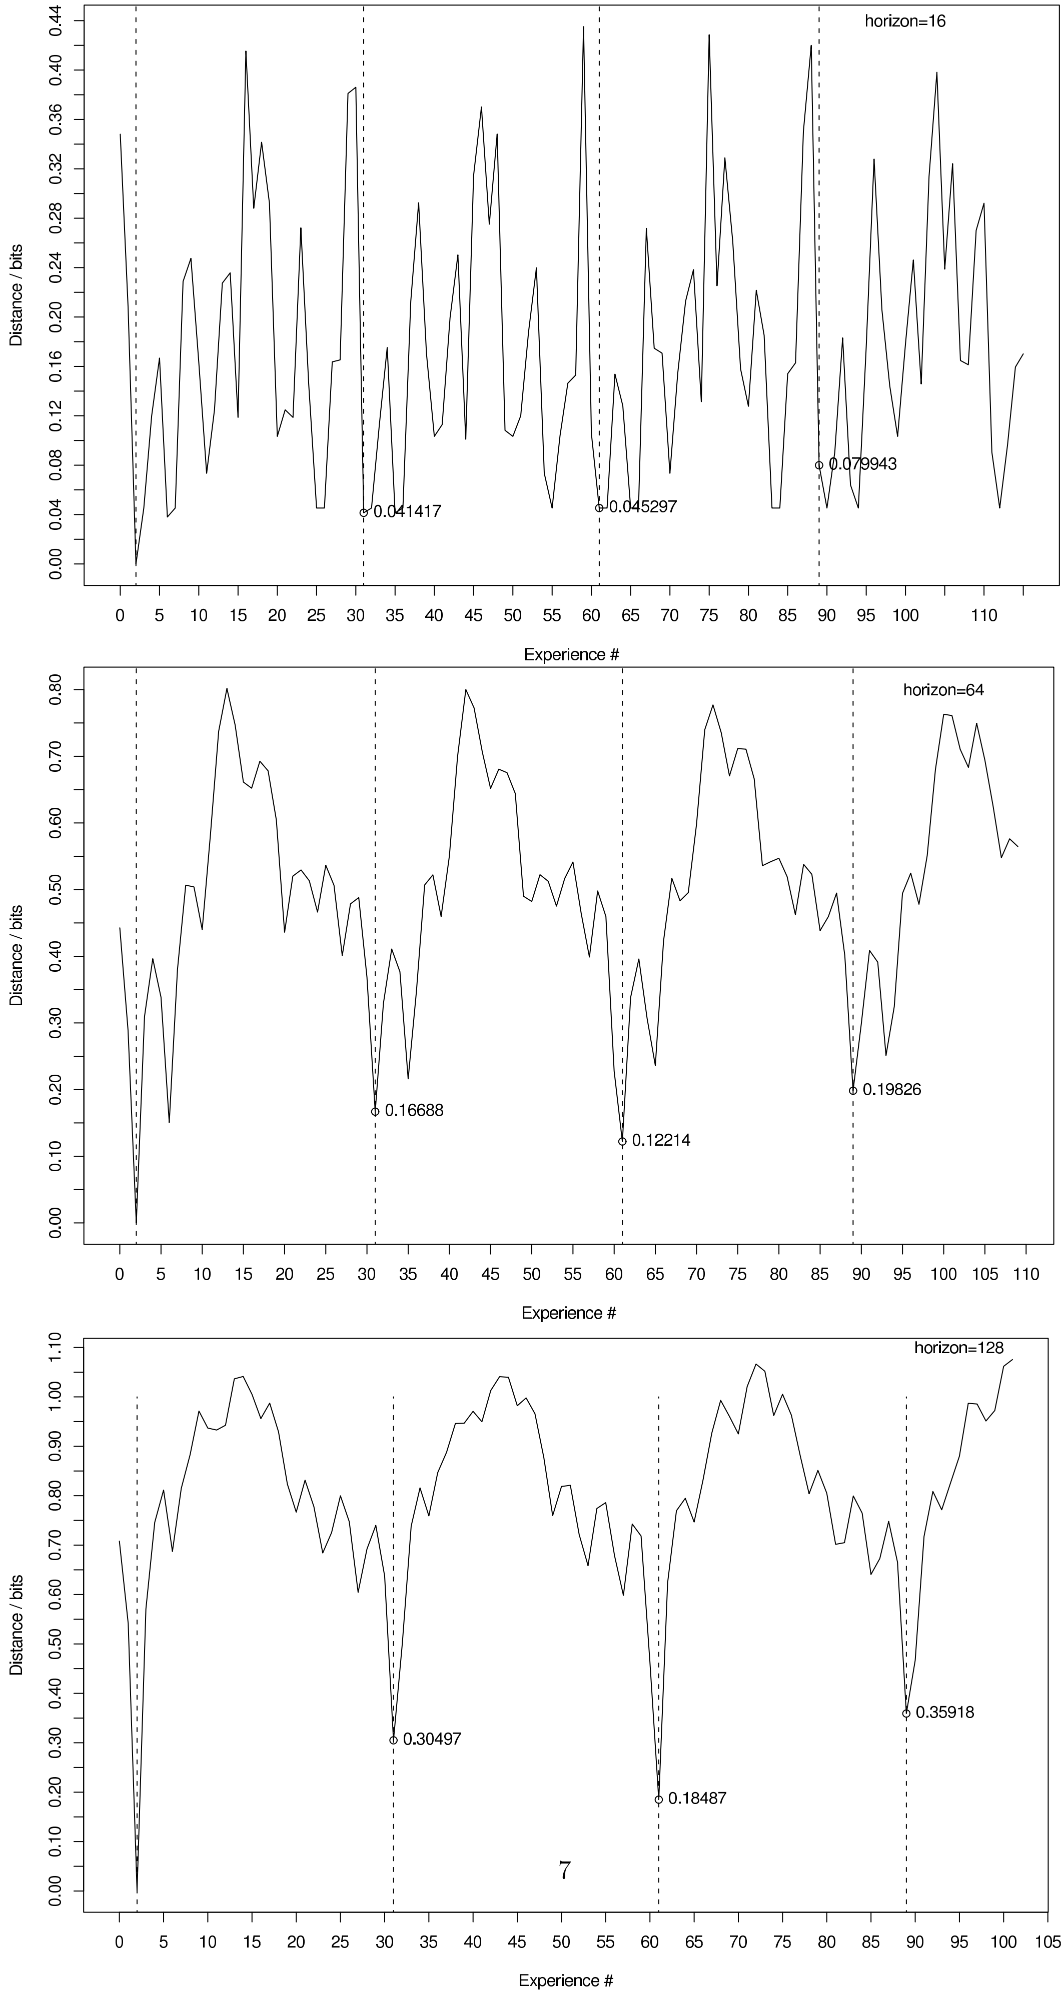

Figure 5: Experience distances for 3 different horizons (h=16, 64, 128). Horizon 16 (top) is not long enough to include the light in the history.

SOCIAL NEUROSCIENCE, 2006, 1 (2), 77–89

Psychology Press
Taylor & Francis Group

Language in shadow

Luciano Fadiga
University of Ferrara, Ferrara, Italy and Italian Institute of Technology, Genova, Italy

Laila Craighero, Maddalena Fabbri Destro, and Livio Finos
University of Ferrara, Ferrara, Italy

Nathalie Cotillon-Williams and Andrew T. Smith
Royal Holloway, University of London, London, UK

Umberto Castiello
Royal Holloway, University of London, London, UK and University of Padova, Padova, Italy

The recent finding that Broca's area, the motor center for speech, is activated during action observation lends support to the idea that human language may have evolved from neural substrates already involved in gesture recognition. Although fascinating, this hypothesis can be questioned because while observing actions of others we may evoke some internal, verbal description of the observed scene. Here we present fMRI evidence that the involvement of Broca's area during action observation is genuine. Observation of meaningful hand shadows resembling moving animals induces a bilateral activation of frontal language areas. This activation survives the subtraction of activation by semantically equivalent stimuli, as well as by meaningless hand movements. Our results demonstrate that Broca's area plays a role in interpreting actions of others. It might act as a motor-assembly system, which links and interprets motor sequences for both speech and hand gestures.

INTRODUCTION

Several theories have been proposed to explain the origins of human language. They can be grouped into two main categories. According to "classical" theories, language is a peculiarly human ability based on a neural substrate newly developed for the purpose (Chomsky, 1966; Pinker, 1994). Theories of the second "evolutionary" category consider human language as the evolutionary refinement of an implicit communication system, already present in lower primates, based on a set of hand/mouth goal-directed action representations (Armstrong, Stokoe, & Wilcox, 1995; Corballis, 2002; Rizzolatti & Arbib, 1998). The classical view is supported by the existence in humans of a cortical network of areas that become active during verbal communication. This network includes the temporal–parietal junction (Wernicke's area), commonly thought to be involved in sensory processing of speech, and the inferior frontal gyrus (Broca's area) classically considered to be the speech motor center.

Correspondence should be addressed to: Luciano Fadiga, Department of Biomedical Sciences and Advanced Therapies, Section of Human Physiology, University of Ferrara, via Fossato di Mortara 17/19, I-44100 Ferrara, Italy. E-mail: fdl@unife.it

This work was supported by EC grants ROBOT-CUB, NEUROBOTICS and CONTACT to LF and LC by European Science Foundation (OMLL Eurocores) and by Italian Ministry of Education grants to LF, and by a Research Strategy Fund to UC.

We thank Attilio Mina for helpful bibliographic suggestions on animal hand shadows, Carlo Truzzi for creating the hand shadows stimuli, and Rosario Canto for technical support.

www.psypress.com/socialneuroscience DOI:10.1080/17470910600976430

However, the existence of areas exclusively devoted to language has increasingly been challenged by experimental evidence showing that Broca's area and its homologue in the right hemisphere also become active during the observation of hand/mouth actions performed by other individuals (Aziz-Zadeh, Koski, Zaidel, Mazziotta, & Iacoboni, 2006; Decety et al., 1997; Decety & Chaminade, 2003; Grafton, Arbib, Fadiga, & Rizzolatti, 1996; Grèzes, Costes, & Decety, 1998; Grèzes, Armony, Rowe, & Passingham, 2003; Iacoboni, Woods, Brass, Bekkering, Mazziotta, & Rizzolatti, 1999; Rizzolatti et al., 1996). This finding seems to favor the evolutionary hypothesis, supporting the idea that the linguistic processing that characterizes Broca's area may be closely related to gesture processing. Moreover, comparative cytoarchitectonic studies have shown a similarity between human Broca's area and monkey area F5 (Matelli, Luppino, & Rizzolatti, 1985; Petrides, 2006; Petrides & Pandya, 1997; von Bonin & Bailey, 1947), a premotor cortical region that contains neurons discharging both when the monkey acts on objects and when it sees similar actions being made by other individuals (Di Pellegrino, Fadiga, Fogassi, Gallese, & Rizzolatti, 1992; Gallese, Fadiga, Fogassi, & Rizzolatti, 1996; Rizzolatti, Fadiga, Gallese, & Fogassi, 1996). These neurons, called ''mirror neurons,'' may allow the understanding of actions made by others and might provide a neural substrate for an implicit communication system in animals. It has been proposed that this primitive ''gestural'' communication may be at the root of the evolution of human language (Rizzolatti & Arbib, 1998). Although fascinating, this theoretical framework can be challenged by invoking an alternative interpretation, more conservative and fitting with classical theories of the origin of language. According to this view, humans are automatically compelled to covertly verbalize what they observe. The involvement of Broca's area during action observation may therefore reflect inner, sub-vocal speech generation (Grèzes & Decety, 2001). Clearly, an experiment determining which of these two interpretations is correct, could provide a fundamental insight into the origins of language.

We investigated here the possibility that Broca's area and its homologue in the right hemisphere become specifically active during the observation of a particular category of hand gestures: hand shadows representing animals opening their mouths. These stimuli have been selected for two main reasons. First, by using these stimuli it was possible to design an fMRI experiment in which any activation due to covert verbalization could be removed by subtraction: the activation evoked while observing videos representing stimuli belonging to the same semantic set (i.e., real animals opening their mouths), and expected to elicit similar covert verbalization, could be subtracted from activity elicited by hand shadows of animals opening their mouths. Residual activation in Broca's area after this critical subtraction would demonstrate the involvement of ''speech-related'' frontal areas in processing meaningful hand gestures. Second, hand shadows only implicitly ''contain'' the hand creating them. Thus they are interesting stimuli that might be used to answer the question of how detailed a hand gesture must be in order to activate the mirror-neuron system. The results we present here support the idea that Broca's area is specifically involved during meaningful action observation and that this activation is independent of any internal verbal description of the seen scene. Moreover, they demonstrate that the mirror-neuron system becomes active even if the pictorial details of the moving hand are not explicitly visible. In the case of our stimuli, the brain ''sees'' the performing hand also behind the appearance.

METHODS

Participants were 10 healthy volunteers (6 females and 4 males; age range 19–32, mean 23). All had normal or corrected vision, no past neurological or psychiatric history and no structural brain abnormality. Informed consent was obtained according to procedures approved by the Royal Holloway Ethics Committee. Throughout the experiment, subjects performed the same task, which was to carefully observe the stimuli, which were back projected onto a screen visible through a mirror mounted on the MRI head coil (visual angle, $15^\circ \times 20^\circ$ approximately). Stimuli were of six types: (1) movies of actual human hands performing meaningless movements; (2) movies of the shadows of human hands representing animals opening their mouths; (3) movies of real animals opening their mouths, plus, as controls, three further movies representing a sequence of still images taken from the previously described three videos. Hand movements were performed by a professional shadow artist. All stimuli were enclosed in a rectangular frame (Figure 1 and online supplementary materials), in a 640×480 pixel array and were shown in grey scale. Each movie lasted 15 seconds, and contained a sequence of 7 different moving/static stimuli (e.g., dog, cow, pig, bird, etc., all opening their mouths). The experiment was conducted as a series of scanning sessions, each lasting 4 minutes. Each session contained eight blocks. In each session two different movie types were presented in an alternated order (see Figure 1, bottom). Each subject completed six sessions. The order of sessions was varied randomly across subjects. The six sessions contrasted the following pairs of movie types: (1) C1 = moving animal hand shadows, C2 = static animal hand shadows; (2) C1 = moving real hands, C2 = static real hands; (3) C1 = moving real animals, C2 = static real animals; (4) C1 = moving animal hand shadows, C2 = moving real animals; (5) C1 = moving animal hand shadows, C2 = moving real hands; (6) C1 = moving real hands, C2 = moving real animals. Whole-brain fMRI data were acquired on a 3T scanner (Siemens Trio) equipped with an RF volume headcoil. Functional images were obtained with a gradient echo-planar T2* sequence using blood oxygenation level-dependent (BOLD) contrast, each comprising a full-brain volume of 48 contiguous axial slices (3 mm thickness, 3×3 mm in-plane voxel size). Volumes were acquired continuously with a repetition time (TR) of 3 seconds. A total of 80 scans were acquired for each participant in a single session (4 minutes), with the first 2 volumes subsequently discarded to allow for T1 equilibration effects. Functional MRI data were analyzed using statistical parametric mapping software (SPM2, Wellcome Department of Cognitive Neurology, London). Individual scans were realigned, spatially normalized and transformed into a standard stereotaxic space, and spatially smoothed by a 6 mm FWHM Gaussian kernel, using standard SPM methods. A high-pass temporal filter (cut-off 120 seconds) was applied to the time series. Considering the relatively low number of participants, a high-threshold, corrected, fixed effects analysis, was first performed by each experimental condition. Pixels were identified as significantly activated if $p < .001$ (FDR corrected for multiple comparisons) and the cluster size exceeded 20 voxels. The activated voxels surviving this procedure were superimposed on the standard SPM2 inflated brain (Figure 1). Clusters of activation were anatomically characterized according to their centers of mass activity with the aid of Talairach co-ordinates (Talairach & Tournoux, 1988), of the Muenster T2T converter (http://neurologie.uni-muenster.de/T2T/t2tconv/conv3d.html) and by taking into account the prominent sulcal landmarks. Furthermore, as far as Broca's region is concerned, a hypothesis-driven analysis was performed for sessions (3)–(6). In this analysis, a more restrictive statistical criterion was used (group analysis on individual subjects analysis, small volume correction approach, cluster size >20 voxels). Only significant voxels ($p < .005$) within the most permissive border of cytoarchitectonically defined probability maps (Amunts, Schleicher, Burgel, Mohlberg, Uylings, & Zilles, 1999) were considered. This last analysis was performed with the aid of the Anatomy SPM toolbox (Eickhoff et al., 2005). Subjects' lips were video-monitored during the whole scanning procedure. No speech-related muscle activity was detectable during video presentation. The absence of speech-related motor activity during video presentation was assessed in a pilot experiment, on a set of different subjects, looking at the same videos presented in the scanner while electromyography of tongue muscles was recorded according to the technique used by Fadiga, Craighero, Buccino, and Rizzolatti (2002).

RESULTS

During the fMRI scanning volunteers observed videos representing: (1) the shadows of human hands depicting animals opening and closing their mouths; (2) human hands executing sequences of meaningless finger movements; or (3) real animals opening their mouths. Brain activations were compared between pairs of conditions in a block design. In addition (4, 5, 6), each condition was contrasted with a "static" condition, in which the same stimuli presented in the movie were shown as static pictures (e.g., stills of animals presented for the same time as the corresponding videos). The comparison between the first three "moving" conditions with each corresponding "static" one, controls for nonspecific activations and emphasizes the action component of the gesture. Figure 1 shows, superimposed, the results of the moving vs. static contrasts for animal hand shadows and real animals conditions (red and green spots, respectively). In addition to largely overlapping occipito-parietal activations, a specific differential activation emerged in the anterior

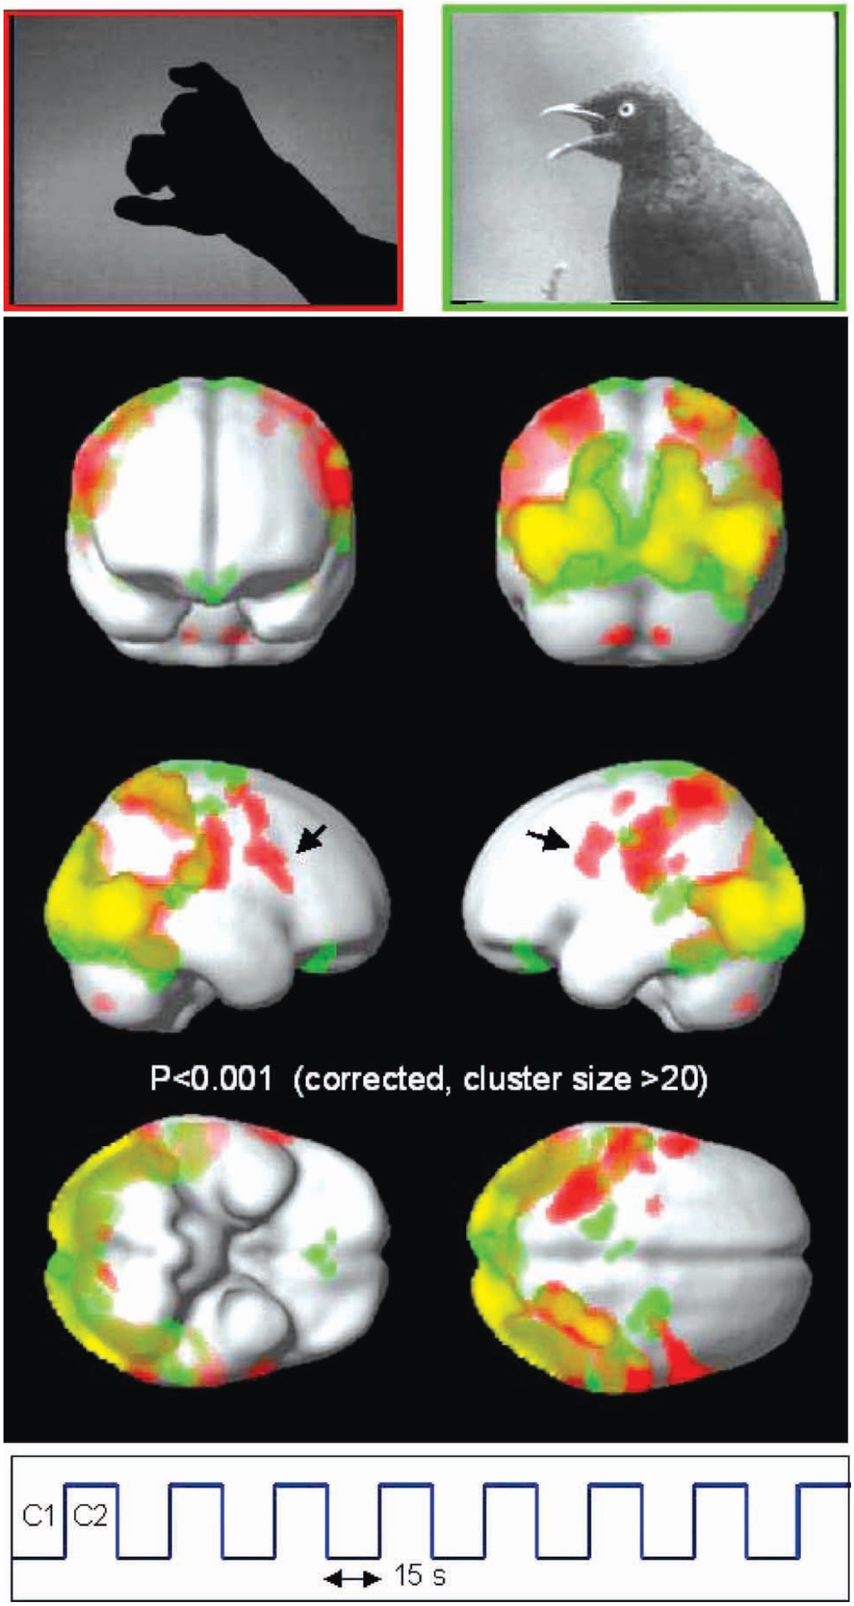

Figure 1 (*See opposite for caption*)

part of the brain. Animal hand shadows strongly activated left parietal cortex, pre- and post-central gyri (bilaterally), and, more interestingly for the purpose of the present study, bilateral inferior frontal gyrus (BA 44 and 45). Conversely, the only frontal activation reaching significance in the moving vs. static contrast for real animals was located in bilateral BA 6, close to the premotor activation shown in an fMRI experiment by Buccino et al. (2004) when subjects observed mouth actions performed by monkeys and dogs. This location may therefore correspond to a premotor region where a species-independent mirror-neuron system for mouth actions is present in humans. A discussion of non-frontal activations is beyond the scope of the present paper, however the above threshold activation foci are listed in Table 1. The results shown in Figure 1 on one side seem to rule out the possibility that the inferior frontal activity induced by action viewing is due to covert speech, on the other side indicate that the shadows of animals opening their mouths, although clearly depicting animals and not hands, convey implicit information about the human being moving her hand in creating them. Indeed, they evoke an activation pattern superimposable on that evoked by hand action observation (Buccino et al., 2001; Grafton et al., 1996; Grèzes et al., 2003; see Figure 1 and Table 1). To interpret this result, it may be important to stress the peculiar nature of hand shadows: although they are created by moving hands, the professionalism of the artist creating them is such that the hand is never obvious. Nevertheless, the mirror-neuron system is activated. The possibility we favor is that the brain "sees" the hand behind the shadow. This possibility is supported by recent data demonstrating that monkey mirror neurons become active even if the final part of the grasping movement is performed behind a screen (Umiltà et al., 2001). Consequently, the human mirror system (or at least part of it) seems to act more as an active interpreter than as a passive perceiver.

The bilateral activation of inferior frontal gyrus shown in Figure 1 during observation of animal hand shadows cannot yet be attributed to covert verbalization. This is because it survives the subtraction of still images representing the same stimuli presented in the moving condition, which might also evoke internal naming. It could be argued, however, that videos of moving animals and animal shadows are dynamic and richer in details than their static controls, and might more powerfully evoke a semantic representation of the observed scene, but this cannot be stated with confidence. We therefore made a direct comparison between moving animal hand shadows and moving real animals. We narrowed the region of interest from the whole brain (as in Figure 1) to bilateral BA 44, the main target of our study. This hypothesis-driven analysis was performed by looking at voxels within the most permissive borders of the probabilistic map of this area provided by Amunts et al. (1999) by taking as significance threshold the p value of .005 (random effect analysis). The results of this comparison are shown in Figure 2B. As already suggested by Figure 1 and Table 1, right and (more interestingly) left frontal clusters survived this subtraction. The first one was located in right BA 44 ($X=58$, $Y=12$, $Z=24$), an area known to be involved during observation of biological action, either meaningful or meaningless (Grèzes et al., 1998; Iacoboni et al., 1999). The second one is symmetrically positioned on the left side ($X=-50$, $Y=4$, $Z=22$). Finally, two additional clusters were present in Broca's region. One was more posterior, in that part of the inferior frontal gyrus classically considered as speech related ($X=-58$, $Y=12$, $Z=14$; *pars opercularis*) and one more anterior, within area 45 according to the Talairach and Tournoux atlas (1988), ($X=-45$, $Y=32$, $Z=14$; *pars triangularis*). This finding agrees with our hypothesis and demonstrates that the activation of Broca's area during action understanding is independent of internal verbalization: if an individual is compelled to verbalize internally when a hand-shadow representing an animal is presented, the same individual should also verbalize during the observation of real animals.

The finding that Broca's area involvement during observation of hand shadows is not

Figure 1 (opposite). Cortical activation pattern during observation of animal hand shadows and real animals. Significantly activated voxels ($p<.001$, fixed effects analysis) in the moving animal shadows and moving real animals conditions after subtraction of the static controls. Activity related to animal shadows (red clusters) is superimposed on that from real animals (green clusters). Those brain regions activated during both tasks are depicted in yellow. In the lowermost part of the figure the experimental time-course for each contrast is shown (i.e., C1, moving; C2, static). Note the almost complete absence of frontal activation for real animals in comparison to animal shadows, which bilaterally activate the inferior frontal gyrus (arrows).

TABLE 1

Montreal Neurological Institute (MNI) and Talairach (TAL) co-ordinates and *T*-values of the foci activated during observation of moving animal hand shadows, real hands, and real animals, after subtraction of static conditions

	MNI			*TAL*			
	x	*y*	*z*	*x*	*y*	*z*	*T-value*
Animal hand shadows							
Inferior frontal gyrus							
BA 44							
R	62	8	24	61	9	22	4.71
L	−62	8	24	−61	9	22	4.58
BA 45							
R	58	22	14	57	22	12	3.37
Precentral gyrus							
BA 6							
R	54	0	36	53	2	33	5.29
	60	2	34	59	4	31	4.90
L	−60	−2	36	−59	0	33	5.09
	−60	0	18	−59	1	17	3.56
	−60	−12	38	−59	−10	36	5.24
BA 4							
R	54	−18	36	53	−16	34	7.97
Postcentral gyrus							
BA 40							
L	−58	−22	16	−57	−21	16	5.92
BA 3							
L	−32	−38	52	−32	−34	50	6.53
BA 2							
R	32	−40	66	32	−36	63	5.18
Superior parietal lobule							
BA 7							
R	26	−48	64	26	−44	61	6.24
Cuneus							
BA 18							
R	16	−100	6	16	−97	10	10.43
	20	−96	8	20	−93	12	12.36
L	−16	−104	2	−16	−101	7	5.09
Middle occipital gyrus							
BA 18							
R	42	−90	8	42	−87	12	7.78
L	−28	−96	2	−28	−93	6	6.95
BA 19							
L	−40	−86	0	−40	−83	4	7.28
Insula							
BA 13							
R	54	−40	20	54	−38	20	6.35
Middle temporal gyrus							
BA 37							
R	54	−70	2	53	−68	5	10.95
Temporal fusiform gyrus							
BA 37							
R	42	−44	−16	42	−43	−11	7.52
L	−44	−44	−16	−44	−43	−11	6.44
Cerebellum							
R	10	−80	−44	8	−80	−42	6.53
L	−8	−78	−44	−8	−77	−33	5.61
Amygdala							
R	22	−2	−22	22	−3	−18	3.85
L	−18	−2	−22	−18	−3	−18	4.15

TABLE 1 (*Continued*)

	MNI			*TAL*			
	x	*y*	*z*	*x*	*y*	*z*	*T-value*
Real hands							
Inferior frontal gyrus							
BA 44							
R	62	16	28	61	17	25	5.67
L	−60	10	26	−60	12	23	3.64
Middle frontal gyrus							
BA 6							
R	34	−6	62	34	−3	57	5.30
L	−24	−8	52	−24	−5	48	4.18
Superior frontal gyrus							
BA 10							
R	4	58	28	4	58	23	6.22
L	−26	54	−2	−27	52	−4	4.14
Postcentral gyrus							
BA 3							
R	52	−20	40	51	−18	38	4.71
BA 7							
R	28	−50	66	28	−46	63	12.50
BA 40							
R	46	−32	54	46	−29	51	4.58
L	−66	−22	14	−65	−21	14	3.37
Inferior parietal lobule							
BA 40							
R	62	−30	28	61	−28	27	4.02
	32	−42	58	32	−38	55	8.23
L	−50	−30	32	−50	−28	31	5.09
	−34	−42	58	−34	−38	55	6.51
Superior parietal lobule							
BA 7							
L	−32	−56	62	−33	−51	60	4.94
Cuneus							
BA 18							
L	−20	−94	6	−20	−91	10	7.13
Middle occipital gyrus							
BA 19							
L	−54	−76	2	−53	−74	6	7.42
	−36	−62	14	−36	−59	16	7.60
Lingual gyrus							
BA 17							
R	10	−96	−8	10	−93	−2	5.74
Inferior occipital gyrus							
BA 18							
R	30	−94	−12	30	−92	−6	6.10
L	−26	−96	−8	−26	−93	−2	6.69
Middle temporal gyrus							
BA 37							
R	52	−68	2	51	−66	5	7.78
Temporal fusiform gyrus							
BA 37							
R	46	−42	−18	46	−41	−13	5.45
L	−46	−44	−16	−46	−43	−11	4.59
Cerebellum							
R	20	−82	−28	20	−81	−20	4.11
L	−28	−56	−50	−28	−56	−39	5.10
Parahippocampal gyrus							
BA 34							
R	30	6	−18	30	5	−15	6.75
Insula							
R	52	−22	18	51	−20	18	6.05

TABLE 1 (*Continued*)

	MNI			*TAL*			
	x	*y*	*z*	*x*	*y*	*z*	*T-value*
Globus pallidus							
R	16	2	2	16	2	2	4.10
L	−16	0	−2	−16	0	−2	3.59
Real animals							
Precentral gyrus							
BA 4							
L	−62	−14	38	−61	−12	36	6.61
BA 6							
R	36	−2	34	36	0	31	5.93
	32	−12	52	32	−9	48	4.48
L	−30	0	38	−30	2	35	10.43
Middle frontal gyrus							
BA 47							
L	−46	48	−2	−46	46	−4	3.92
Superior frontal gyrus							
BA 6							
R	30	−8	72	30	−4	67	3.66
BA 8							
R	20	30	56	20	32	50	6.13
Postcentral gyrus							
BA 2							
L	−38	−40	70	−38	−36	66	4.56
BA 3							
R	34	−36	54	34	−32	51	3.26
L	−18	−44	76	−18	−39	72	3.35
Precuneus							
BA 7							
L	−8	−54	46	−8	−50	45	6.08
Middle occipital gyrus							
BA 18							
R	20	−98	14	20	−94	18	11.75
L	−20	−94	10	−20	−91	14	14.91
BA 19							
R	38	−80	2	38	−77	6	7.21
L	−40	−86	−2	−40	−83	2	9.88
	−54	−76	2	−53	−74	6	8.43
Lingual gyrus							
BA 18							
R	14	−86	−14	14	−84	−8	5.40
L	−6	−88	−10	−6	−86	−4	5.24
Limbic lobe-uncus							
BA 28							
L	−28	8	−24	−28	7	−21	7.00
Superior temporal gyrus							
BA 41							
L	−46	−40	12	−46	−38	13	3.77
Middle temporal gyrus							
BA 21							
L	−54	−24	−8	−53	−24	−6	3.54
BA 37							
R	52	−70	4	51	−68	7	8.03
Temporal fusiform gyrus							
BA 37							
R	44	−42	−20	44	−41	−15	6.50
L	−46	−50	−20	−46	−49	−14	4.66
Cerebellum							
R	46	−58	−32	46	−58	−24	3.55
L	−48	−58	−32	−48	−58	−24	3.78

TABLE 1 (*Continued*)

	MNI			*TAL*			
	x	*y*	*z*	*x*	*y*	*z*	*T-value*
Amygdala							
L	−34	−6	−16	−34	−6	−13	3.70
Globus Pallidus							
L	−24	−10	−4	−24	−10	−3	4.93
Anterior Cingulate BA 24							
L	−4	36	8	−4	35	6	4.70

Note: BA = Brodmann area; R = right hemisphere; L = left hemisphere; x, y, z = co-ordinates.

explainable in terms of internal speech suggests that, in agreement with other data in the literature, this area may play a crucial role in action understanding. However, it remains the fact that Broca's area is known as a speech area and its involvement during overt and covert speech production has clearly been demonstrated recently (Palmer, Rosen, Ojemann, Buckner, Kelley, & Petersen, 2001). The hypothesis we favor is that this area participates in verbal communication because it represents the product of the evolutionary development of a precursor already present in monkeys: the mirror neurons area F5, that portion of the ventral premotor cortex where hand/mouth actions are represented (see Petrides, 2006). Accordingly, in agreement with the characteristics of area F5 neurons, Broca's area should respond much better to goal-directed action than to simple, meaningless movements. To test this possibility, we performed two further comparisons: (1) The observation of human hands performing meaningless finger movements versus the observation of moving real animals opening their mouths, to determine how much of the *pars opercularis* activation was due to the observation of meaningless hand movements; (2) The observation of animal hand shadows versus the observation of meaningless finger movements, to pit the presence of hands against the presence of meaning. The results are shown in Figure 2, C and D, respectively. After comparison (1), although the more dorsal, bilateral, BA 44 activation was still present (Figure 2C, left: X = −56, Y = 10, Z = 26; right: X = 58, Y = 10, Z = 24), no voxels above significance were located in the *pars opercularis* of Broca's area. This demonstrates that finger movements per se do not activate specifically that part of Broca's area. In contrast, after comparison (2) a significant activation was present in the left *pars opercularis* (Figure 2D; X = −58, Y = 6, Z = 4), demonstrating the involvement of Broca's area *pars opercularis* in processing actions of others, particularly when meaningful and thus, implicitly, communicative.

DISCUSSION

The finding that animal hand shadows but not real animals or meaningless finger movements activate that part of Broca's region most intimately involved in verbal communication support a similarity between these stimuli and spoken words. Animal hand shadows are formed by meaningless finger movements combined to evoke a meaning in the observer through the shape appearing on a screen. Thus, when one looks at them, the representation of an animal opening its mouth is evoked. Words that form sentences are formed by individually meaningless movements (phonoarticulatory acts), which appropriately combined and segmented convey meanings and representations. Does this twofold involvement of Broca's area reflect a specific role played by it in decoding actions and particularly communicative ones? A positive answer to this question arises, in our view, from the finding that when observation of meaningless finger movements is subtracted from observation of animal hand-shadows, an activation of the left *pars opercularis* persists.

The activation of Broca's area during gestural communication has already been shown in deaf signers, both during production and perception. This was interpreted as a vicariant involvement of Broca's area because of its verbal specialization (Horwitz et al., 2003). In other terms, according to this interpretation, Broca's area is activated because, by signing, deaf people express linguistic concepts. In our study

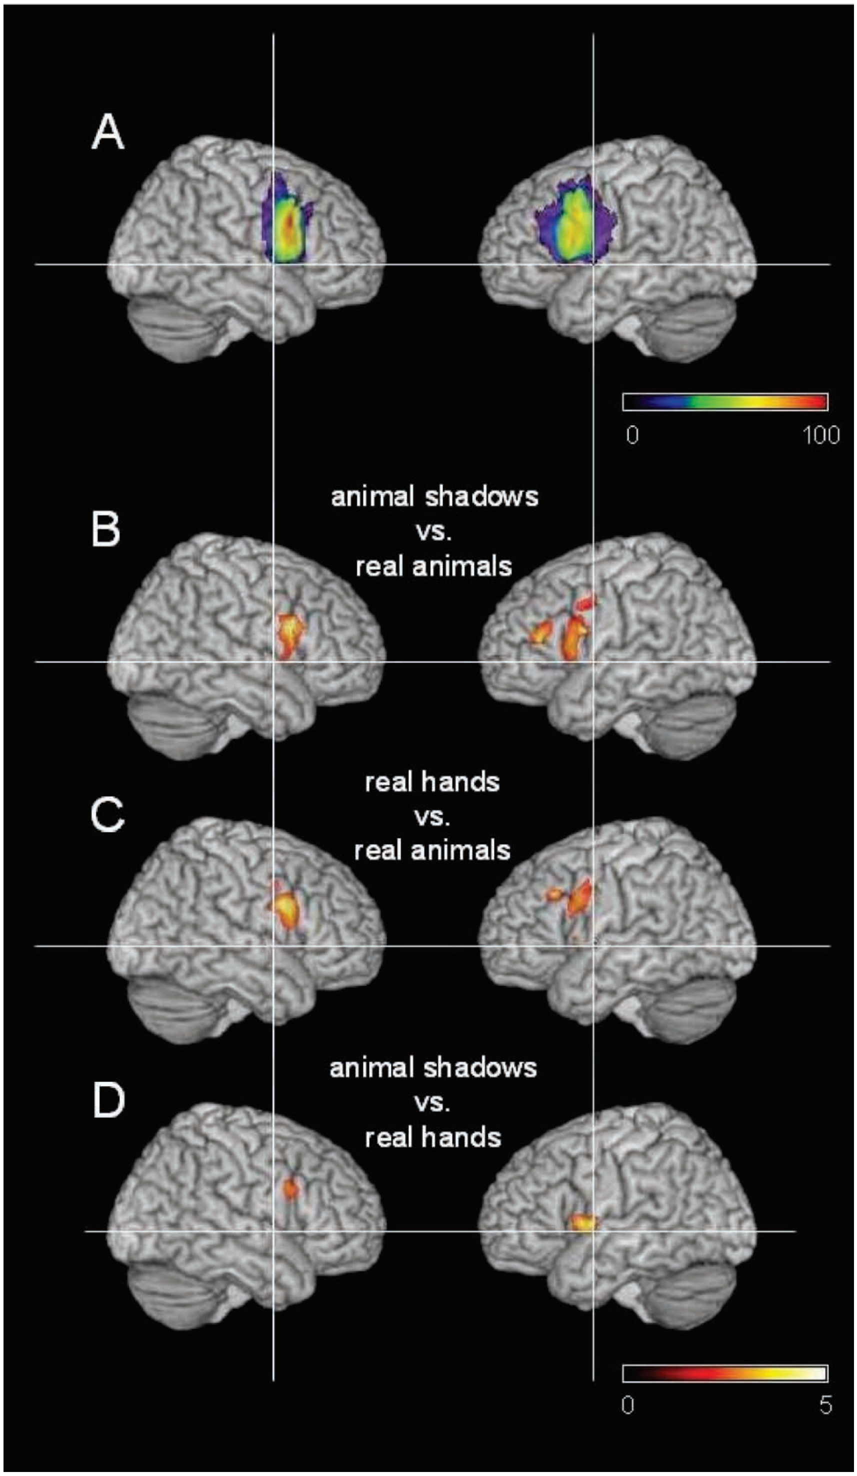

Figure 2 (*See opposite for caption*)

participants were presented with communicative hand gestures but, in contrast to studies investigating deaf people, the gestures were non-symbolic even if able to address in an unambiguous way a specific concept (e.g., a barking dog). Thus, we show here, the involvement of Broca's region can not be explained in terms of linguistic decoding of the gesture meaning. Conversely, the results indicate that Broca's region is involved in the understanding of communicative gestures.

How can this ''perceptual'' function be reconciled with the universally accepted *motor* role of Broca's area for speech? One possible interpretation is that Broca's area, due to its premotor origin, is involved in the assembly of meaningless sequence of action units (whether finger or phonoarticulatory movements) into meaningful representations. This elaboration process may proceed in two directions. In production, Broca's area recruits movement units to generate words/hand actions. In perception, Broca's area, being the human homologue of monkey area F5, addresses the vocabulary of speech/hand actions, which form the template for action recognition. Our hypothesis is that, in origin, Broca's area precursor was involved in generating/extracting action meanings by organizing/interpreting motor sequences in terms of goal. Subsequently, this ability might have been generalized during the evolution that gave this area the capability to deal with meanings (and rules), which share similar hierarchical and sequential structures with the motor system (Fadiga, Craighero, & Roy, 2006).

This proposal is in agreement with fMRI investigations that indicate that Broca's area is not always activated during speech listening. In a recent experiment Wilson, Saygin, Sereno, and Iacoboni (2004) carried out an fMRI study in which subjects (1) passively listened to monosyllables and (2) produced the same speech sounds. Results showed a substantial bilateral overlap between regions activated during the two conditions, mainly in the superior part of ventral premotor cortex. Conversely, the activation of Broca's region was present only in some of the studied subjects, in our view because the task did not require any meaning extraction. This interpretation is in line with brain imaging studies indicating that, in speech comprehension, Broca's area is mainly activated during processing of syntactic aspects (Bookheimer, 2002). Luria (1966) had already noticed that Broca's area patients made comprehension errors in syntactically complex sentences such as passive constructions. Finally, data coming from cortical stimulation of collaborating patients undergoing neurosurgery, showed that the electrical stimulation of the Broca's area produced comprehension deficits, particularly evident in the case of ''complex auditory verbal instructions and visual semantic material'' (Schaffler, Luders, Dinner, Lesser, & Chelune, 1993). The data of the present experiment, together with the series of evidence presented above, are in agreement with those theories on the origins of human language that consider it as the evolutionary refinement of an implicit communication system based on hand/mouth goal-directed action representations (Armstrong et al., 1995; Corballis, 2002; Rizzolatti & Arbib, 1998). This possibility finds further support from a recent experiment based on the analysis of brain MRIs of three great ape species (*Pan troglodytes*, *Pan paniscus* and *Gorilla gorilla*) showing that the extension of BA 44 is larger in the left hemisphere than in the right. While a similar asymmetry in humans has been correlated with language dominance (Cantalupo & Hopkins, 2001), this hypothesis does not fit in the case of apes. It might be, however, indicative of an initial specialization of BA 44 for communication. In fact, in captive great apes manual gestures are both referential and intentional, and are preferentially produced by the right hand. Moreover, this right-hand bias is consistently greater when gesturing is accompanied by vocalization (Hopkins & Leavens, 1998).

In conclusion, our results support a common origin for human speech and gestural communication in non-human primates. It has been proposed that the development of human speech is a consequence of the fact that the precursor of Broca's area was endowed, before the emergence of speech, with a gestural recognition system (Rizzolatti & Arbib, 1998). Here we have taken a step forward, empirically showing for the first time that human Broca's area is not an exclusive ''speech'' center but, most probably, a motor assembly center in which communicative

Figure 2 (opposite). Results of the analysis focused on bilateral area 44. (A) Cytoarchitectonically defined probability map of the location of left and right area 44, drawn on the Colin27T1 standard brain on the basis of Juelich-MNI database (Amunts et al., 1999). The white cross superimposed on each brain indicates the origin of the co-ordinates system ($x = y = z = 0$). The correspondence between colors and percent probability is given by the upper color bar. (B), (C) and (D), significant voxels ($p < .005$, random effects analysis) falling inside area 44, as defined by the probability map shown in (A), in the three contrasts indicated in the Figure. Color bar: T-values. Note the similar pattern of right hemisphere activation in (B) and (C), the similar location of the posterior-dorsal activation of the left hemisphere in (B) and (C), and the two additional foci in the *pars opercularis* and *pars triangularis* of Broca's area in (B). Note, in (D), the survival of the activation in *pars opercularis*, after subtraction of real hands from animal hand shadows. When the reverse contrasts were tested (real animals vs. either animal shadows or real hands), the results failed to show any significant activation within area 44.

gestures, whether linguistic or otherwise, are assembled and decoded (Fadiga et al., 2006). It still remains unclear whether hand/speech motor representations are mapped in this area according to a somatotopic organization, or if Broca's area works in a supramodal way, by dealing with effector-independent motor rules.

Manuscript received 31 May 2006
Manuscript accepted 1 August 2006
First published online 6 October 2006

REFERENCES

Amunts, K., Schleicher, A., Burgel, U., Mohlberg, H., Uylings, H. B., & Zilles, K. (1999). Broca's region revisited: Cytoarchitecture and intersubject variability. *Journal of Comparative Neurology*, *412*, 319–341.

Armstrong, A. C., Stokoe, W. C., & Wilcox, S. E. (1995). *Gesture and the nature of language*. Cambridge, UK: Cambridge University Press.

Aziz-Zadeh, L., Koski, L., Zaidel, E., Mazziotta, J., & Iacoboni, M. (2006). Lateralization of the human mirror neuron system. *Journal of Neuroscience*, *26*(11), 2964–2970.

Bookheimer, S. (2002). Functional MRI of language: New approaches to understanding the cortical organization of semantic processing. *Annual Review of Neuroscience*, *25*, 151–188.

Buccino, G., Binkofski, F., Fink, G. R., Fadiga, L., Fogassi, L., Gallese, V., et al. (2001). Action observation activates premotor and parietal areas in a somatotopic manner: An fMRI study. *European Journal of Neuroscience*, *13*, 400–404.

Buccino, G., Lui, F., Canessa, N., Patteri, I., Lagravinese, G., Benuzzi, F., et al. (2004). Neural circuits involved in the recognition of actions performed by nonconspecifics: An fMRI study. *Journal of Cognitive Neuroscience*, *16*, 114–126.

Cantalupo, C., & Hopkins, W. D. (2001). Asymmetric Broca's area in great apes. *Nature*, *414*(6863), 505.

Chomsky, N. (1966). *Cartesian linguistics*. New York: Harper & Row.

Corballis, M. C. (2002). *From hand to mouth. The origins of language*. Princeton, NJ: Princeton University Press.

Decety, J., & Chaminade, T. (2003). Neural correlates of feeling sympathy. *Neuropsychologia*, *41*, 127–138.

Decety, J., Grèzes, J., Costes, N., Perani, D., Jeannerod, M., Procyk, E., et al. (1997). Brain activity during observation of actions: Influence of action content and subject's strategy. *Brain*, *120*, 1763–1777.

Di Pellegrino, G., Fadiga, L., Fogassi, L., Gallese, V., & Rizzolatti, G. (1992). Understanding motor events: A neurophysiological study. *Experimental Brain Research*, *91*, 176–180.

Eickhoff, S. B., Stephan, K. E., Mohlberg, H., Grefkes, C., Fink, G. R., Amunts, K., et al. (2005). A new SPM toolbox for combining probabilistic cytoarchitectonic maps and functional imaging data. *Neuroimage*, *25*, 1325–1335.

Fadiga, L., Craighero, L., & Roy, A. C. (2006). Broca's area: A speech area? In Y. Grodzinsky & K. Amunts (Eds.), *Broca's region*. New York: Oxford University Press.

Fadiga, L., Craighero, L., Buccino, G., & Rizzolatti, G. (2002). Speech listening specifically modulates the excitability of tongue muscles: A TMS study. *European Journal of Neuroscience*, *15*, 399–402.

Gallese, V., Fadiga, L., Fogassi, L., & Rizzolatti, G. (1996). Action recognition in the premotor cortex. *Brain*, *119*, 593–609.

Grafton, S. T., Arbib, M. A., Fadiga, L., & Rizzolatti, G. (1996). Localization of grasp representations in humans by PET: 2. Observation compared with imagination. *Experimental Brain Research*, *112*, 103–111.

Grèzes, J., & Decety, J. (2001). Functional anatomy of execution, mental simulation, observation, and verb generation of actions: A meta-analysis. *Human Brain Mapping*, *12*, 1–19.

Grèzes, 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.

Grèzes, J., Costes, N., & Decety, J. (1998). Top-down effect of strategy on the perception of human biological motion: A PET investigation. *Cognitive Neuropsychology*, *15*, 553–582.

Hopkins, W. D., & Leavens, D. A. (1998). Hand use and gestural communication in chimpanzees (*Pan troglodytes*). *Journal of Comparative Psychology*, *112*(1), 95–99.

Horwitz, B., Amunts, K., Bhattacharyya, R., Patkin, D., Jeffries, K., Zilles, K., et al. (2003). Activation of Broca's area during the production of spoken and signed language: A combined cytoarchitectonic mapping and PET analysis. *Neuropsychologia*, *41*, 1868–1876.

Iacoboni, M., Woods, R. P., Brass, M., Bekkering, H., Mazziotta, J. C., & Rizzolatti, G. (1999). Cortical mechanisms of human imitation. *Science*, *286*, 2526–2528.

Luria, A. (1966). *The higher cortical function in man*. New York: Basic Books.

Matelli, M., Luppino, G., & Rizzolatti, G. (1985). Patterns of cytochrome oxidase activity in the frontal agranular cortex of macaque monkey. *Behavioral Brain Research*, *18*, 125–137.

Palmer, E. D., Rosen, H. J., Ojemann, J. G., Buckner, R. L., Kelley, W. M., & Petersen, S. E. (2001). An event-related fMRI study of overt and covert word stem completion. *Neuroimage*, *14*, 182–193.

Petrides, M. (2006). Broca's area in the human and nonhuman primate brain. In Y. Grodzinsky & K. Amunts (Eds.), *Broca's region*. New York: Oxford University Press.

Petrides, M., & Pandya, D. N. (1997). Comparative architectonic analysis of the human and the macaque frontal cortex. In F. Boller & J. Grafman (Eds.),

Handbook of neuropsychology (Vol. IX, pp. 17–58). New York: Elsevier.
Pinker, S. (1994). *The language instinct: How the mind creates language.* New York: William Morrow.
Rizzolatti, G., & Arbib, M. A. (1998). Language within our grasp. *Trends in Neuroscience*, *21*, 188–194.
Rizzolatti, G., Fadiga, L., Gallese, V., & Fogassi, L. (1996). Premotor cortex and the recognition of motor actions. *Cognitive Brain Research*, *3*, 131–141.
Rizzolatti, G., Fadiga, L., Matelli, M., Bettinardi, V., Paulesu, E., Perani, D., et al. (1996). Localization of grasp representations in human by PET: 1. Observation versus execution. *Experimental Brain Research*, *111*, 246–252.
Schaffler, L., Luders, H. O., Dinner, D. S., Lesser, R. P., & Chelune, G. J. (1993). Comprehension deficits elicited by electrical stimulation of Broca's area. *Brain*, *116*, 695–715.
Talairach, J., & Tournoux, P. (1988). *Co-planar stereotactic atlas of the human brain*. New York: Georg Thieme Verlag.
Umiltà, M. A., Kohler, E., Gallese, V., Fogassi, L., Fadiga, L., Keysers, C., et al. (2001). I know what you are doing: A neurophysiological study. *Neuron*, *31*, 155–165.
Von Bonin, G., & Bailey, P. (1947). *The neocortex of Macaca mulatta*. Urbana, IL: University of Illinois Press.
Wilson, S. M., Saygin, A. P., Sereno, M. I., & Iacoboni, M. (2004). Listening to speech activates motor areas involved in speech production. *Nature Neuroscience*, *7*(7), 701–702.