ROBotic Open-architecture Technology for
Cognition, Understanding and Behavior

Project No. 004370

RobotCub
Development of a Cognitive Humanoid Cub

Instrument: Integrated Project
Thematic Priority: IST - Cognitive Systems

D6.1 Results from Computational/Robotic Models of Gesture Communication
WP6 - Gesture Communication

Due date: **01/08/2006**
Submission Date: **24/09/2006**

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: **1.0**

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

Contents

1 Introduction: Before Gesture and Communication

1.1 Executive Summary

This document is the second deliverable for D6.1, comprising a summary of the research work in grounding gesture communication and results from robotic models and experiments. We report on ongoing research in development and use of gestural communication using interaction histories in social games, perception of robot expressions using KASPAR an interactive robot developed at the University of Hertfordshire, adaptation of a robot to detected styles of play, and work to develop an eye-contact detection module for the iCub.

We present results from a study of interaction dynamics between a human and a robot that develops the capacity to play the early learning game "peekaboo" using an underlying interaction history based control architecture. A robotic model was developed that uses a history of interactions to direct future actions extending earlier work in this work package. The interaction history is realised as a metric space of sensorimotor experiences based on the informational relationships between experiences. "Peekaboo" is important as it is considered to contribute developmentally to infant understanding and practice of social interaction, providing the scaffolding upon which infants can co-regulate their emotional expressions with others, build social expectations and establish primary intersubjectivity [28]. Study of the rules, timing of gestures, and the ontogeny of such games in a artificial agents will enhance our understanding of the mechanisms of the development of early gestural communication capabilities.

We report on research in how the interaction between a human and robot can be affected by the design of the face and in particular the perception of certain robotic expressions (in particular the smile). This is important in laying the basis for communicative interactions using facial expressions and gestures. KASPAR, the expressive, interactive humanoid robot platform for studying human robot interaction and kinesics developed at Hertfordshire, is used in these trials as well as in the child-robot interaction studies reported in deliverable 5.4, and we report here on the progress in its development.

We present a proof-of-concept of a robot adapting its behaviour on-line during interactions with a human according to detected play styles. Play is important as a vehicle for learning and developing skills in a variety of areas including communication and social skills. In the context of autonomous robots using play in development, it is useful for the robot to be able to engage in different styles of play both within a single developmental stage and between stages.

1.2 Relation to Project Objectives

Deliverable 6.1 and work conducted in WP6 contributes primarily to PO-3 (as stated in Annex 1, Section 2), that is "the study and implementation of early stages of human cognitive development in an embodied artificial system". Specifically, it focuses on parts *iv* and *v*, that is, "learning and regulating interaction dynamics" and "developing autobiographic memory based on interaction histories".
The further development of the dynamically constructed, embodied interaction history grounded in the sensorimotor experience of the robot, with extension to predicting and directing future action contributes to the construction of a shared autobiographic memory and interaction history as detailed in PO-3-*v*. Further, the study of the "peekaboo" gestural communication scenario contributes to the

learning of interaction dynamics as stated in PO-3-*iv*.
In relation to the specific project objectives (Annex 1, Section 8) the work in this deliverable contributes to SO-3 (b) "the ability of understanding and exploiting simple gestures to interact socially".

1.3 Approach and Motivation

We take a bottom-up approach to studying gesture and communication, starting from signalling and interaction at a local level and exploring emerging interactional structure. Our study is based on a model of cognition as the development and activity of an embodied dynamical system structurally coupled to its environment, which develops in sophistication in response to a history of interactions with the environment (including the social environment). Cognitive structures arise from the recurrent sensorimotor patterns that enable and scaffold increasingly complex perceptually guided interaction.

Trevarthen [31] describes the importance of rhythm and timing and inter-subjectivity in early communicative interactions of infants with a caregiver, and terms these *protoconversations*. Motivated by this viewpoint we propose looking at how a robotic agent can develop capabilities to engage in these kinds of protoconversations with a human partner, and choose early interaction games such as "peekaboo" as vehicles for this study.
In terms of mechanisms by which communicative interaction and supporting cognitive processes can come about, we are motivated in particular by embodied dynamical systems and their development [30, 8, 14]. Non-linear dynamical systems show stable states, attractors and transitions between them. The activity and stability of a dynamical system is governed by system parameters and changes in these parameters cause transitions between attractors. Cognitive processes and in particular remembering can be understood in these terms with the attractors of the system being memories, recall of which is caused by (sensory or internal) inputs to the system causing transitions between stable attractors and states.
A characteristic of the cognitive systems of human infants is that they do not base actions only on immediate sensor input. Instead, as for post-reactive systems in general, they are able to act on a wider time-horizon [23] basing actions on a history of previous interactions. Furthermore they base actions on predictions of sensory inputs and states of the world, and this extension of the time-horizon into the future goes from sub-second predictive movement of muscles through expectations of future world and personal states to types of planning and goal directed action.
Our approach recognises that this wide temporal horizon is grounded in the sensorimotor experience of the agent, and this motivates research into defining a behavioural model which has at its core a dynamically created and recreated interaction history that emerges from the sensorimotor flow of information. New experience is compared using information theoretic concepts to known experience and used to guide and predict future behaviour and interaction. Representation and significance of events within this history are emergent quantities arising from interaction with the environment and natural reward systems.
An important aspect of our approach is for any history of sensory input, actions and interactions to be without externally imposed representational constraints. This is motivated by the desire to ground any internal structure in the sensorimotor coupling with the environment and to circumvent the symbol-grounding problem[1]. From the perspective of cognition as the development and activity of an embodied dynamical system, the experiential history consists of time-extended structures in

[1]The *symbol grounding problem* refers to the difficulty in connecting the semantics of arbitrary symbols to real world entities or events [1], this is not to say that there are no symbols, only that they are sufficiently well grounded.

the sensorimotor space and the space of experiences (described below) is organised on the direct sensorimotor history of the agent.

2 Research Areas: Interaction Histories

2.1 Peekaboo Early Interaction Game

The development of gestural communication skills is grounded in the early interaction games that infants play. Therefore in the study of the ontogeny of social interaction, gestural communication and turn-taking in artificial agents, it is instructive to look at the kinds of interactions 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 care-giver 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[2] 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 [21, 32].
Bruner and Sherwood [5] studied peekaboo from the viewpoint of play and learning of the rules and structures of games. They also recognise 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 largely non-verbal communication. The hiding of the face is one such gesture, and the vocalisation, and the showing of pleasure (laughing) are others. In order for the protoconversation to proceed successfully, the gestures must be made by either party at the times expected by the players, and the absence or mistiming 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 vocalisation on renewed contact.
In all this, the rhythm and timing of the interaction are crucial and, Bruner and Sherwood suggest 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 [26, pp. 424-425].
In relation to the development of social cognition in infants, “peekaboo” and other social interaction games that are characterised by a building and then releasing of tension in cyclic phases are important as they are considered to contribute developmentally to infant understanding and practice 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 [28].

2.2 Interaction History Architecture

We describe an architecture where an embodied robotic agent can make use of an interaction history to guide ontological development to act appropriately in a changing environment. The direct senso-

[2] In the context of humour, 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 [32].

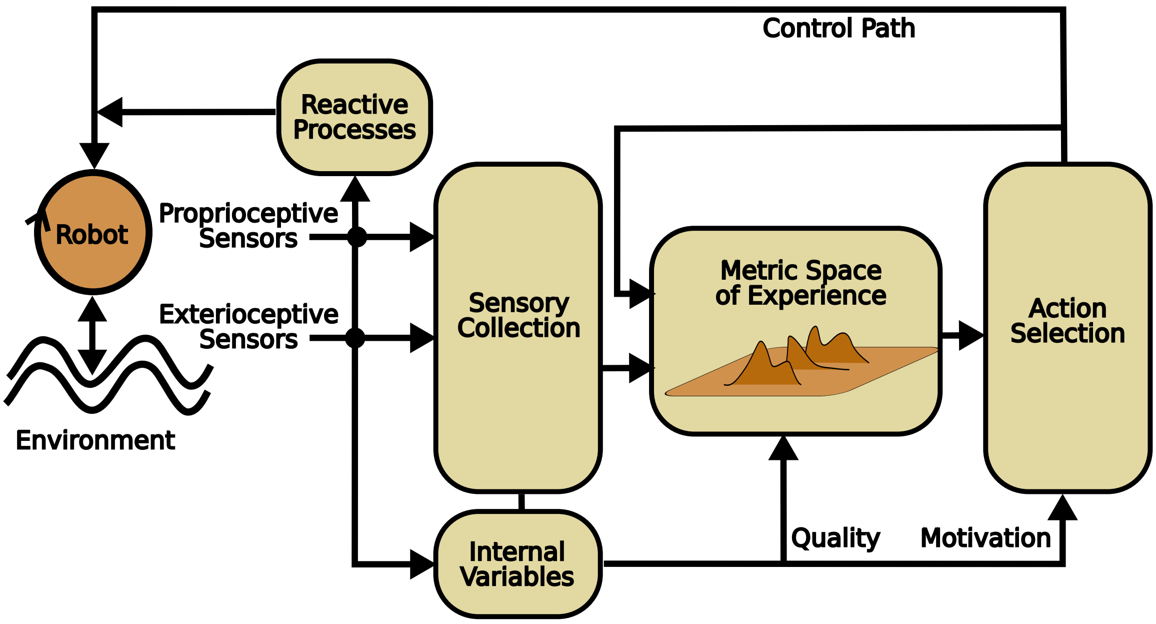

Figure 1: Schematic of the Interaction history based control architecture.

rimotor motor history of the agent is used to create grounded experiences of different lengths which can be compared with one another using a metric measure based on the information distance between them. The agent acts on the basis of its experiences and the choice of action depends, in part, on the feedback of reward from the environment. This architecture was initially explored in [19], where the efficacy of the experience space was verified by the robot's use of the history to predict the future position of a ball.
To relate experiences with other experiences in an interaction history, we use information distance measures [29, 6] and a mathematical concept of experience and the relations between them. These are defined in [24] (See Appendix A). Information distance related techniques have been successfully used in the past, for instance, to compare behaviours from the perspective of the agent [17, 13] and for an agent to infer a model of its own sensory and actuator apparatus by acting in the environment [25]. This suggests that behaviour can be guided by moving in a continually constructed space of experiences by selecting appropriate actions that will move the agent closer to particular experiences, return to familiar ones and explore new ones [18].

We describe a computational model (Figure 1) that demonstrates how such interaction histories can be explicitly integrated into the control of a robot. The basic architecture consists of processes to acquire sensory and motor data from the robot as it acts in the environment. From this a metric space consisting of past interaction experiences is constructed by comparing new experiences to previous ones in terms of their information distance. A process then selects past experiences near (i.e. with low information distance) to the current experience. The selection is also based on the values of internal variables that change according to a motivational system. The action following the chosen past experience becomes the next action of the agent. Finally, there is an internal feedback process that adjusts the values of internal variables associated with any experience when it has been used to select future action, making it more or less likely to be chosen in the future.

For a more detailed description of the interaction history architecture see published papers [19, 20] reproduced in Appendix B1.1 and B1.2.

2.3 Experiments with the Interaction History Architecture

Experiments were conducted using an implementation of the Interaction History Architecture realised in a Sony Aibo robot. The robot interacted in controlled conditions with both a human subject or a fixed image of a human face. An initial experiment established the efficacy of the metric-space of experience using a simple task of predicting the path of a ball by matching the relevant experience from the history of experiences. A second preliminary study assessed the architecture as a whole in controlling a robot in playing peekaboo with a human partner. A third experiment investigated the effect of the horizon length of experiences on the ability to play the game and was conducted using a static image of a face for the robot to interact with.
In these experiments the role of the vocalisations as encouragement is replaced by an internally generated dynamic system of variables altering in response to key stimuli, i.e. presentation of a human face.

2.3.1 Experiment 1: Metric-space and ball prediction

Given a robot acting in an environment, how well can it predict future events based on its recent history of experience?
In this experiment the metric space of experience was tested in absence of the action control loop. The experimental condition was that the robot was stationary, and a pink ball was moved in the air in view of the robot's visual system. The path of the ball in each trial included repeated vertical, horizontal and circular movements. The robot was equipped with software that could localise the position of the centre of the ball with respect to the video frame.
The sensorimotor experience of the robot used to create experiences consisted of all sensor values (including IR distance and buttons) and motor positions. Visual sensory flow was derived by extracting average red colour values in 9 regions of the video frame. Additionally the horizontal and vertical location of the ball constituted two further sensors.
The time for the ball to describe a circle (or to move horizontally or vertically for a complete cycle) was approximately 6000ms, and the horizon length of the experiences in the experiment was approximately 4800ms or 40 timesteps of 120ms duration. Thus the horizon length was shorter than, but on the same scale as a single cycle of the repeated behaviour and the experiences would comprise $\frac{4}{5}$ of a cycle.
It was found that the experience matching was successful in that the matched experience corresponded well with the current experience such that, projecting forward, the subsequent ball position could be correctly determined. For full results and discussions see [19] reproduced as Appendix B1.1.

2.3.2 Experiment 2: Peekaboo with human partner

This second experiment utilised the full interaction history architecture to control the actions of the robot by referring to a developing space of experience from which to select actions. The scenario was the peekaboo game described where the robot played the role of the child in the later version of the game where the child hides its face. The game was simplified to not include an audio component which, along with the emotional dimensions of expression generation and recognition has the function of providing encouraging motivation to ensure continuation and reward for participating. This

was replaced by an internally generated motivational subsystem that was driven by coupled sets of equations controlled by the presence or absence of a face in the view of the robot's camera. The motivational subsystem had the characteristic that the presence of a face alone would not result in a very high values for the motivational variables but repeated disappearance and reappearance of a face would result in increasingly higher values.
The experiences stored in the history were tagged with the value of the motivational variables and this, in combination with the proximity of experiences to the target (current) experience, provided a selection criteria for a similar high valued experience. The action taken following the selected experience was chosen as the next action for the robot. Where a suitable experience could not be found (as was more likely in the early interactions) then a random choice of action was taken. This is similar to a form of body babbling [16] and serves to explore the possibilities offered by interaction with the environment.
The results from this trial showed that it was possible for the robot to choose a series of actions that resulted in peekaboo-type behaviour (i.e. the robot repeatedly hiding and revealing its face) when faced with a human partner. The experience space was observed to have been modified enhancing certain specific experiences. Again for full results and discussions see [19], Appendix B1.1.

2.3.3 Experiment 3: Effect of horizon length

The purpose of this investigation was initially to evaluate whether the model for development based on interaction history performed better than random action selection for the task of playing the game of peekaboo. Secondly, the hypothesis that the horizon length of experience would affect the ability to develop the capacity to play the peekaboo game 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.
This experiment again utilised the simplified peekaboo scenario with the internal motivational system replacing the audio and emotive components of the interaction, however in this case the human partner was also replaced with a static image of a face.
Six trials of 2 minute duration for each horizon length of 8, 16, 32, 64 and 128 timesteps (0.96, 1.92, 3.84, 7.68 and 15.36 seconds respectively) were run. For comparison, a further 6 trials were run where the action selection was random and not based on history. In each of the trials the experiential metric space started unpopulated.
Key results were firstly that despite the random action selection resulting in short periods of peekaboo-type behaviour, that longer sustained peekaboo-type behaviour only occurred in the experience driven trials. Furthermore there were specific situations in the experience driven trials where an interrupted peekaboo sequence was recommenced on seeing a face, whereas for the random trial, no clear relationship between seeing a face and playing of peekaboo was seen. A second important result was that the horizon length of experience that results in the peekaboo cyclic behaviour was indeed of the same order as that of the length of cycle itself. This would indicate a need for an experience space that was built up of many different horizon lengths, thus covering different periods of cyclic or recurring behaviour, and for methods for adaptive selection of appropriate temporal horizons.
The experimental results are published as [20] and are reproduced as Appendix B1.2 to this report.

3 Research Areas: Interactive Humanoid

3.1 Introduction

As robots enter everyday life and start to interact with ordinary people, questions surrounding their appearance [10], expressions [4], and of the regulation of interaction dynamics between human and robot are becoming increasingly important. Our perception of a robot can be strongly influenced by its general appearance and in particular its facial appearance.
We conducted research with the goal of advising on robot design for the purposes of human-robot interaction (HRI) in developmental robotics, and synthesising relevant ideas from narrative art design, the psychology of face recognition, and recent HRI studies into robot faces. We discuss the effects of the uncanny valley and the use of iconicity and its relationship to the self-other perceptive divide, as well as abstractness and realism, while classifying existing designs along these dimensions.
KASPAR, a minimally expressive HRI research robot being developed at the University of Hertfordshire provides a test-bed for some of these concepts and we discuss the current state of development. The first study using KASPAR looked at human perceptions of robot expressions (smiles) was conducted and is discussed below.
The robot KASPAR has also been used to support a series of "Wizard-of-Oz" type of experiments where KASPAR interacts via non-verbal communication with typically developing children and children with autism. The results of trials studying the kinesics, or rhythm and timing of non-verbal gesture and communication, are reported in deliverable D5.4

3.2 KASPAR - Progress Report

During RobotCub meetings and discussions in 2005 it became apparent that the issue of facial expressions had not been addressed specifically in the original workplan, in particular the use of an articulated face and the effect of expressions as a feedback device for people interacting with the iCub had not been considered. To investigate these issues as part of task 6.2 (early communication behaviours) it was agreed by the RobotCub coordinator that University of Hertfordshire will pursue a small scale pilot study into minimal expressiveness in a humanoid robot head, due to the group's expertise in social robotics, human-robot interaction as well as affective computing and expression through interaction in robots.
As a result, Hertfordshire developed KASPAR during the second half of the first year, a low-cost child-sized humanoid HRI research robot to perform research into fundamental communication behaviour.
In the second year, control boards and camera capture devices were integrated into the robot. Software was developed to control the low-level function of the motors and capture of video images. Further API software was developed to allow the robot to be controlled using a Wizard-of-Oz paradigm, providing the operator the ability to programme and recall gestures and assign them to keystrokes, and also for later autonomous functioning.
The hardware was enhanced to include two articulated 3 DOF arms with fixed (moulded) hands. This allowed the robot to execute a wider range of gestures and carry out tasks such as playing a drum.
Further development is planned to include simple hands with cable driven digits. There will be independent control of the thumb and forefinger with the distal 3 fingers tied and controlled together. Software development planned includes wrapping the control software in YARP2 for the purpose of compatibility with RobotCub software.

3.3 Design of Robot Faces

The effect of the aesthetic design of a robot is an area of scientific study that has often been neglected in the past. A notable exception is the 'uncanny valley' proposed by Masahiro Mori [22, 9] suggesting that the acceptance of a humanoid robot increases as realism increases, but with a catastrophic fall in acceptance just as the robot approaches human-like realism. In more recent work DiSalvo an colleagues [10] proposed that human robot design should balance considerations of 'human-ness', 'robot-ness' and 'product-ness' to manage the user acceptance of the robot.

In consideration of the human face, studies would suggest that symmetry, youthfulness and skin condition are important factors in perception of attractiveness [12] (although there are claims to the contrary, see [27]). In his book *Understanding Comics* [15], Scott McCloud introduces a triangular design space for cartoon faces, placing faces in a space where the extremes represent realistic, iconic and abstract faces. Dautenhahn suggests that the more iconic a face appears, the more people it can represent and so can aid believability [7].

This design space can be used for interactive and developmental robot designs, and may prove informative in choosing appropriate designs for different niches of robots. For instance, if a robot is required to carry out tasks on a person's behalf, then an iconic design may possibly allow the user to project themselves onto the robot and gain acceptance of the robot's role. This is further discussed in [3] and [2] (included as Appendix B2.1).

3.4 Experiment - Perception of Robot Smiles

A study was conducted to investigate people's perception of the expressions made by KASPAR, and in particular the smile. This would provide a necessary baseline to inform future experiments where KASPAR's expressions will be used in interaction dynamics experiments.

The experiment presented subjects with videos of KASPAR smiling and asked them to rate the smile for perceived "happiness" of the robot, and the how "appealing" the smile looked. Parameters that were manipulated were the 'size' of the smile and the transition time from a neutral expression to the smile, with a static image of the various smiles and the neutral expression providing further conditions. The hypotheses were: that static expressions would appear less appealing than dynamic transitions, that natural transition speeds would be more attractive than abrupt transitions, and that the larger the smile, the better would be the recognition of happiness in the expression. The results supported all three hypotheses, with the finding that, although static expressions appear less appealing than natural speed dynamic transitions as expected, abrupt transition were perceived as even less appealing.

The full results and discussions of this experiment are in [2] which is reproduced in the appendices.

4 Research Areas: Detecting and Adapting to Different Styles of Play

Play is important as a vehicle for learning and developing skills in a variety of areas including communication and social skills. In the context of autonomous robots using play in development, it is useful for the robot to be able to engage in different styles of play both within a single developmental stage and between stages. An initial challenge is for the robot to detect different styles of play as they are occurring and adapt its style of play to the circumstances.

We present a proof-of-concept of a robot adapting its behaviour on-line during interactions with a human according to detected play styles. The generic styles of play explored in the study were "gentle" and "strong" and involved a human interacting with the Sony Aibo using touch sensors on the head,

chin, and back of the robot. Sensor data was pre-processed by first quantizing and summing over all input sensors and then using Fast Fourier Transforms (FFT) to extract shift-invariant frequency domain information. The magnitude of the vectors from the FFT were fed into an artificial neural network self-organising map (SOM).
The SOM was trained offline and nodes classified as "gentle" and "strong" according to their activation in response to the type of interaction in the training set. Experiments were then conducted which initially validated that the algorithm running on-line was able to correctly and efficiently classify types of play, and moved on to test the ability for the robot to adapt its type of play on-line in response to the detected type of play. A further experiment investigated the model's response to the type of repetitive play sometimes practised by children with autism.

The full details of the experiments and results as well as discussions on achievements and limitations of the model can be found in [11], reproduced as Appendix B3.1 to this report. This work is conducted in conjunction with the Aurora project (http://homepages.feis.herts.ac.uk/ comqbr/aurora/index.html) which investigates how robots may be used to help children with autism to overcome some of their impairments in social interactions.

5 Research Areas: Gestures EPFL

6 Research Areas: Neuroscience

References

[1] Ronald C. Arkin. *Behavior-based Robotics*. MIT Press, Cambridge, MA, USA, 1998.

[2] Mike Blow, Kerstin Dautenhahn, Andrew Appleby, Chrystopher L. Nehaniv, and David Lee. Perception of robot smiles and dimensions for human-robot interaction design. In *The 15th IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN 06)*, pages 469–474, Hatfield, UK, 6-8 September 2006.

[3] Mike Blow, Kerstin Dautenhahn, Andrew Appleby, Chrystopher L. Nehaniv, and David Lee. Robot faces and the self/other divide. In *1st ACM Annual Conference on Human-Robot Interaction (HRI2006)*, pages 331–332, Salt Lake City, USA, 2-3 March 2006.

[4] C. L. Breazeal. *Designing sociable robots*. The MIT Press, 2002.

[5] Jerome. S. Bruner and V. Sherwood. Peekaboo and the learning of rule structures. In J.S. Bruner, A. Jolly, and K. Sylva, editors, *Play: Its Role in Development and Evolution*, pages 277–285. New York: Penguin, 1975.

[6] J.P. Crutchfield. Information and its metric. In L. Lam and H.C. Morris, editors, *Nonlinear Structures in Physical Systems - Pattern Formation, Chaos and Waves*, pages 119–130. Springer-Verlag, New York, 1990.

[7] K. Dautenhahn. Design spaces and niche spaces of believable social robots. In *Proc. 2002 IEEE Intl. Workshop Robot and Human Interactive Communication (RO-MAN 2002), 25-27 September, Berlin, Germany*, pages 192–197. IEEE Press, 2002.

[8] K. Dautenhahn and T. Christaller. Remembering, rehearsal and empathy - towards a social and embodied cognitive psychology for artifacts. In Seán Ó Nualláin, Paul McKevitt, and Eoghan Mac Aogáin, editors, *Two Sciences of the Mind: Readings in cognitive science and consciousness*, pages 257–282. John Benjamins North America Inc., 1996.

[9] Kerstin Dautenhahn. Socially intelligent agents in human primate culture. In Robert Trappl and Sabine Payr, editors, *Agent Culture: Human-Agent Interaction in a Multicultural World*, chapter 3, pages 45–71. Lawrence Erlbaum Associates, 2004.

[10] C. DiSalvo, F. Gemperle, J. Forlizzi, and S. Kiesler. All robots are not created equal: The design and perception of humanoid robot heads. In *Proc. Designing Interactive Systems*, pages 321–326, June 2002.

[11] Dorothée François, Daniel Polani, and Kerstin Dautenhahn. On-line behaviour classification and adaptation to human-robot interaction styles. Technical report, University of Hertfordshire, School of Computer Science, College Lane, Hatfield, UK. AL10 9AB, September 2006.

[12] Benedict C Jones, Anthony C Little, D Michael Burt, and David I Perrett. When facial attractiveness is only skin deep. *Perception*, 33(5):569 – 576, 2004.

[13] F. Kaplan and V.V. Hafner. 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)*, pages 129–134. IEEE, 2005.

[14] J. A. S. Kelso and J. P. Scholz. Cooperative phenomena in biological motion. In H. Haken, editor, *Complex systems - Operational approaches in neurobiology, physics and computers*, pages 124–149. Springer-Verlag, Berlin, 1985.

[15] Scott McCloud. *Understanding Comics: The Invisible Art*. Harper Collins Publishers, Inc., 1993.

[16] A. Meltzoff and M. Moore. Explaining facial imitation: a theoretical model. *Early Development and Parenting*, 6:179–192, 1997.

[17] N. A. Mirza, C. L. Nehaniv, K. Dautenhahn, and R. te Boekhorst. Using sensory-motor phase-plots to characterise robot-environment interactions. In *Proc. of 6th IEEE International Symposium on Computational Intelligence in Robotics and Automation (CIRA2005)*, pages 581–586, 2005.

[18] N. A. Mirza, C. L. Nehaniv, K. Dautenhahn, and R. te Boekhorst. Using temporal information distance to locate sensorimotor experience in a metric space. In *Proc. of 2005 IEEE Congress on Evolutionary Computation (CEC2005)*, volume 1, pages 150–157, Edinburgh, Scotland, UK, 2-5 September 2005. IEEE Press.

[19] N. A. Mirza, C. L. Nehaniv, K. Dautenhahn, and R. te Boekhorst. Interaction histories: From experience to action and back again. In *Proceedings of the 5th IEEE International Conference on Development and Learning (ICDL 2006)*, Bloomington, IN, USA, 2006. ISBN 0-9786456-0-X.

[20] N. A. Mirza, C. L. Nehaniv, K. Dautenhahn, and R. te Boekhorst. Peekaboo: Effect of experience length on the interaction history driven ontogeny of a robot. In *Proceedings the of 6th International Conference on Epigenetic Robotics*, pages 71–78, Paris, France, 20-22 September 2006. Lund University Cognitive Studies.

[21] Diane P. F. Montague and Arlene S. Walker-Andrews. Peekaboo: A new look at infants perception of emotion expression. *Developmental Psychology*, 37(6):826–838, 2001.

[22] Masahiro Mori. Bukimi no tani [The uncanny valley]. *Energy*, 7:33–35, 1970. (in Japanese).

[23] C. L. Nehaniv, D. Polani, K. Dautenhahn, R. te Boekhorst, and L. Cañamero. Meaningful information, sensor evolution, and the temporal horizon of embodied organisms. In *Artificial Life VIII*, pages 345–349. MIT Press, 2002.

[24] Chrystopher L. Nehaniv. Sensorimotor experience and its metrics. In *Proc. of 2005 IEEE Congress on Evolutionary Computation*, volume 1, pages 142–149, Edinburgh, Scotland, UK, 2-5 September 2005. IEEE Press.

[25] L. Olsson, C. L. Nehaniv, and D. Polani. From unknown sensors and actuators to visually guided movement. In *Proceedings the of International Conference on Development and Learning (ICDL 2005)*, pages 1–6. IEEE Computer Society Press, 2005.

[26] Roy D. Pea. The social and technological dimensions of scaffolding and related theoretical concepts for learning, education and human activity. *The Journal of the Learning Sciences*, 13(3):423–451, 2004.

[27] D.I. Perrett, K. May, and S. Yoshikawa. Attractive characteristics of female faces: preference for non-average shape. *Nature*, 368:239–242, 1994.

[28] Philippe Rochat, Jane G. Querido, and Tricia Striano. Emerging sensitivity to the timing and structure of protoconversation in early infancy. *Developmental Psychology*, 35(4):950–957, 1999.

[29] Claude E. Shannon. A mathematical theory of communication. *Bell Systems Technical Journal*, 27:379–423 and 623–656, 1948.

[30] E. Thelen and L. B. Smith. *A Dynamic Systems Approach to the Development of Cognition and Action*. MIT Press, 1994.

[31] C. Trevarthen. Musicality and the intrinsic motive pulse: evidence from human psychobiology and infant communication. *Musicae Scientae. Special Issue*, pages 155–215, 1999.

[32] Thomas C. Veatch. A theory of humour. *International Journal of Humour Research*, 11(2):161–175, 1998.

7 Publications

Research activities conducted wholly or partly within WP6 of the RobotCub project have resulted in the following publications in peer-reviewed international conference proceedings between 1st October 2005 and 30th September 2006.

Mike Blow, Kerstin Dautenhahn, Andrew Appleby, Chrystopher L. Nehaniv, and David Lee. Robot faces and the self/other divide, *1st ACM Annual Conference on Human-Robot Interaction (HRI2006)*, pages 331–332, Salt Lake City, USA, 2-3 March 2006.

Mike Blow, Kerstin Dautenhahn, Andrew Appleby, Chrystopher L. Nehaniv, and David Lee. Perception of robot smiles and dimensions for human-robot interaction design, In *The 15th IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN 06)*, pages 469–474, Hatfield, UK, 6-8 September 2006.

N. A. Mirza, C. L. Nehaniv, K. Dautenhahn, and R. te Boekhorst. Interaction histories: From experience to action and back again. In *Proceedings the of 5th International Conference on Development and Learning (ICDL 2006)*, Bloomington, IN, USA, 2006. ISBN 0-9786456-0-X.

N. A. Mirza, C. L. Nehaniv, K. Dautenhahn, and R. te Boekhorst. Peekaboo: Effect of experience length on the interaction history driven ontogeny of a robot. In *Proceedings the of 6th International Conference on Epigenetic Robotics*, pages 71–78, Paris, France, 20-22 September 2006. Lund University Cognitive Studies. 2006.

Dorothée François, Daniel Polani, and Kerstin Dautenhahn. On-line behaviour classification and adaptation to human-robot interaction styles. *Technical Report , University of Hertfordshire, School of Computer Science, College Lane, Hatfield, UK. AL10 9AB*, September 2006.

A Appendix A: Geometry of Experience

A.1 Information Sources as Random Variables

Consider a sensor or effector that can take on various settings or values modeled as a random variable $\mathcal{X}$ changing with time, taking value $x(t) \in X$, where X is the set of its possible values with discrete time t.

A.2 Entropy and Information Distance

Entropy is the information-theoretic measure of uncertainty introduced by Claude Shannon [29] and its units are bits. The *entropy* $H(\mathcal{X})$ of a sensor or actuator $\mathcal{X}$ is then $H(\mathcal{X}) = -\sum_{x \in X} p(x) \log_2 p(x)$, where $p(x)$ gives the probability of value x being taken. *Conditional entropy* $H(\mathcal{X}|\mathcal{Y})$ of a random variable $\mathcal{X}$ given $\mathcal{Y}$ is the amount of uncertainty that remains about the value $\mathcal{X}$ given that the value of $\mathcal{Y}$ is known.

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

where $p(x,y)$ is given by the joint distribution of $\mathcal{X}$ and $\mathcal{Y}$.[3] The *information distance* between $\mathcal{X}$ and $\mathcal{Y}$ is

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

This satisfies the mathematical axioms for a *metric*:

1. $d(\mathcal{X},\mathcal{Y}) = 0$ if and only if $\mathcal{X}$ and $\mathcal{Y}$ are equivalent.[4]
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).

The satisfaction of these axioms is shown by Crutchfield [6]. Thus d defines a geometric structure on any space of jointly distributed information sources such as a robot's sensory, motor and internal variables.

A.3 Sensorimotor Variables with Time Horizons

For a particular agent, in a particular environment, consider a sensorimotor variable $\mathcal{X}$. Its distribution will be affected by the agent-environment interaction. In the context of a particular environment and beginning from a particular moment in time t_0 until a later moment $t_0 + h$ ($h > 0$), we regard the sequence of values $x(t_0), x(t_0+1), \ldots, x(t_0+h-1)$ taken by an information source $\mathcal{X}$ as time-series data from a new random variable $\mathcal{X}_{t_0,h}$, the *sensorimotor variable with temporal horizon h for sensor (or actuator) $\mathcal{X}$ starting at time t_0*, depending on situated experience.

A.4 Information Distance between Time-Shifted Sensorimotor Variables

A particular robot engages in various behaviours and interactions in a particular environment, and we consider two of its sensorimotor variables $\mathcal{X}$ and $\mathcal{Y}$. (Possibly $\mathcal{X} = \mathcal{Y}$.) Consider the values taken

[3] 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.

[4] For information sources, "equivalence" refers to *re-coding equivalence*. That is, the values of $\mathcal{X}$ are a function of those of $\mathcal{Y}$ and vice versa. See [6].

by $\mathcal{X}$ beginning at time t_0 and those of $\mathcal{Y}$ beginning at time t_1. (Possibly $t_0 = t_1$.) Consider the two-component random variable $\mathcal{X}_{t_0,h} \times \mathcal{Y}_{t_1,h}$ with horizon h, whose distribution is estimated from the values $(x(t_0+i), y(t_1+i)) \in X \times Y$. The first component here comes from $\mathcal{X}$, starting from time t_0, and second component comes from $\mathcal{Y}$ with a temporal shift of $t_1 - t_0$ units, starting from time t_1. We can also estimate the probability joint time-shifted distribution and the information distance $d(\mathcal{X}_{t_0,h}, \mathcal{Y}_{t_1,h})$ between $\mathcal{X}$ during the first temporal region and $\mathcal{Y}$ during the second temporal region by measuring the frequencies of occurrence of values (x_{t_0+i}, y_{t_1+i}) as i runs from 0 to $h-1$.
Clearly there are issues related to the size of the temporal horizon h and also the number of values $\mathcal{X}$ and $\mathcal{Y}$ may take that affect the accuracy of these estimates. Also in practice, independent samples of time shifted sensorimotor variables are not available.

A.5 Experience Metric

Consider the set of all sensorimotor variables available to an agent. Suppose there are N such, $\mathcal{X}^1, \ldots, \mathcal{X}^N$. Let $E(t,h) = (\mathcal{X}^1_{t,h}, \ldots, \mathcal{X}^N_{t,h})$ be the (ordered) set of these variables considered over a temporal window of size h starting at t. We call $E(t,h)$ the agent's *experience* from time t having temporal horizon h.
Let $E = E(t,h)$ and $E' = E(t',h)$ be experiences of an agent from time t and t', respectively, both with horizon size h. Define a metric on experiences of temporal horizon h as

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

where d is the information distance.

D is a metric on the set of experiences of an agent having a fixed temporal horizon h and $\overline{D} = \frac{1}{N}D$, the *average experience distance per sensorimotor variable*, is also a metric on the set of experiences of an agent having a fixed temporal horizon h.

The units of D are bits and those of $\overline{D}$ are bits per sensorimotor variable.

Another metric on experience is given by

$$D^*(E,E') = \Sigma_{k=1}^{N} \Sigma_{j=1}^{N} d(\mathcal{X}^k_{t,h}, \mathcal{X}^j_{t',h}).$$

Clearly, $D \leq D^*$.

B Appendix B: Selected Publications

B.1 Interaction Histories

B.1.1 Interaction histories: From experience to action and back again.

N. A. Mirza, C. L. Nehaniv, K. Dautenhahn, and R. te Boekhorst. Interaction histories: From experience to action and back again. In *Proceedings the of 5th International Conference on Development and Learning (ICDL 2006)*, Bloomington, IN, USA, 2006. ISBN 0-9786456-0-X.

Interaction Histories: From Experience to Action and Back Again *

Naeem Assif Mirza, Chrystopher L. Nehaniv, Kerstin Dautenhahn, René te Boekhorst
Adaptive Systems Research Group, School of Computer Science
University of Hertfordshire, College Lane, Hatfield, AL10 9AB. United Kingdom
{N.A.Mirza,C.L.Nehaniv,K.Dautenhahn,R.teBoekhorst}@herts.ac.uk

***Abstract*— We describe an enactive, situated model of interaction history based around a growing, informationally self-structured metric space of experience that is constructed and reconstructed as the robot engages in sensorimotor interactions with objects and people in its environment. The model shows aspects of development and learning through modification of the cognitive structure that forms the basis for action selection as a result of acting in the world. We describe robotic experiments showing prediction of the path of a ball and an interaction game "peekaboo".**

***Index Terms*— Interaction History, Information Theory, Robotic Control Architectures**

I. Introduction

A challenge of research into situated, enactive cognition in robots is to reach beyond reactive architectures to architectures that can reflect the time-extended behaviour characteristic of humans and many animals. We are interested in how cognitive structures in natural and artificial systems can arise that capture the history of interactions and behaviours of an agent actively engaged in its environment, without resorting to symbolic representations of past events.

We introduce an architecture that has at its heart a changing dynamic structure describing the space of experience of the agent or robot. The robot chooses how to behave in the world based on what it has experienced, and this results in further experience and modification of previous experience establishing a tight coupling of experience and action.

This paper proceeds by presents our concept of an interaction history and then describes the model and architecture that we use. Finally we describe experiments conducted on a robot platform that investigate the capabilities of the model.

II. Interaction History

We use a working definition of an *interaction history* as *the temporally extended, dynamically constructed and reconstructed, individual sensorimotor history of an agent situated and acting in its environment including the social environment.* The first key part of the definition is that the agent is situated and actively acting within its environment, that is the history is not a disembodied memory, but an active part of the interaction of the agent and its environment. This follows the idea of structural coupling and enactive cognition of Maturana and Varela [1] and the concept of situated cognition [2]. Remembering is then the effect of historical interactions on the actions of an agent in response to a particular situations [3]. This brings in the next key part of the definition, that the history is dynamically constructed and reconstructed. In other words, interactions with the environment construct the structures that are used for remembering how to act. Thus, memory consists not of static representations of the past that can be recalled with perfect clarity, but rather is the result of an accumulation of interaction with the environment manifesting as current action.

An important aspect of the interaction history is that it is constructed from the perspective of the individual, that is, it is autobiographical in nature. In terms of the accepted separation of memory types due to Endel Tulving [4], this would be 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. However this apparently clear dichotomy is not applicable to a description of interaction history as, through the process of reconstruction, categories and knowledge may emerge from many overlapping experiences, while certain unique events may still stand out and give memory its episodic nature. While we do not claim that an interaction history can describe all aspects of (human) memory, we believe that exploring the features of an interaction history may give insights into the nature of memory as a whole. The final part of the definition that we would highlight is that it need not be representational but must be grounded in the sensorimotor experience of the agent.

A. Extended Temporal Horizon

A robotic agent with an interaction history has the potential to act on an extended temporal horizon [5] resulting in behaviour that goes beyond that of a reactive agent or an affective agent. The distinction is that behaviour will be

*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.

modulated by temporally extended past experience as well as by internal state (affect) and immediately by environmental stimuli (reactivity).

B. Development and Learning

A further aspect of an interaction history which manifests itself as modification of behaviour based on a history of previous interactions is that it can serve to scaffold learning and development of a situated agent. The key here is how previous experience is used to affect current and future behaviour. For example, classical conditioning or a two-process reinforcement learning based on positive and negative reinforcers [6] are potential mechanisms for connecting previous experience with choice of action. Development can be seen as the increasing richness of the connections of experience with action, again mediated by a suitable mechanisms.

III. Enactive Robot Model of Interaction History using Sensorimotor Experience

We describe a computational robotic model (Fig. 1) that illustrates how an interaction history can be integrated into the control of a robot using the concepts described in the previous section.

The basic architecture consists of processes to acquire sensory and motor values from the robot as it acts in the environment, from this a metric space of past interaction experiences is constructed. A further process continuously examines current experience in the context of the space of previous experience and selects actions to execute.

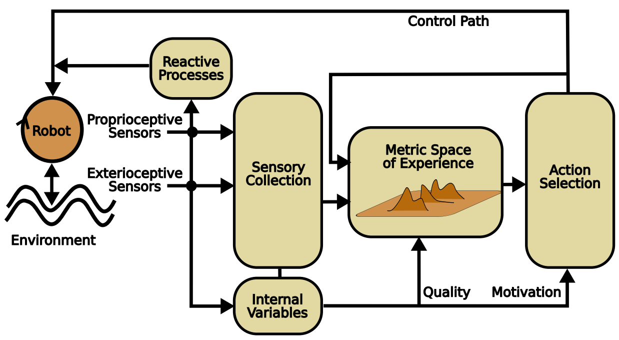

Fig. 1. Interaction history based control architecture.

A. Sensory and Internal Variables

The sensory information available to the robot falls into three broad categories: proprioceptive, exteroceptive and internal. Proprioceptive variables are constructed by sampling motor position and exteroceptive variables are those from sensors such as buttons, infra-red distance, vision[1] and audition[2]. In addition to these, sensory input can also be built from internal variables that might, for instance, indicate drives and motivations, or be the result of processing of raw sensory data e.g. ball position. Sampling is done at regular intervals (between 100-120ms in the experiments presented).

B. Experience Space

The experience space is constructed from overlapping experiences of a particular horizon size with relative positions in the space determined by the informational distance between them (see section *IV*). Many potential experience spaces of different horizon length can be built and co-exist [7].[3].
As the metric landscape of experience is built, each experience is further enhanced with *value attributes* of the experience. These are the instantaneous values of any sensor or internal variable, for example variables indicating "satiation", "battery-level", "contentment" and so forth. Experiences are also annotated with the actions that the robot takes at any timestep (see section *III-C*).

C. Action Selection, Development and Learning

While an experience space can be built without much difficulty, the challenge is how to have experience modulate future action in a meaningful way and to be further shaped by that action. To achieve this goal, a simple mechanism is adopted whereby the robot can execute one of a number of "atomic" actions (or no action) at any timestep[4]. At any timestep the robot can choose an action based on past experience or, if an appropriate one is not found, can choose a random one. The ability to choose a random action has the advantage of emulating body-babbling, i.e. apparently random body movements that have the (hypothesized) purpose of learning the capabilities of the body in an environment [8]. Early in development, there are few experiences on which to draw, so random actions would be chosen more often, and later in development, it is more likely that an appropriate experience (and thus action) can be found. Additionally, with a small probability, the robot may still choose a random action as this may help move out of "local minima", and potentially discover new, more salient experiences.
To choose an action based on experience, the robot first examines the experience landscape for similar experiences near the current one. That is it finds a *candidate experience* with the shortest information distance to the current one. The next action that was executed following that experience is a *candidate action* to be executed next.
The candidate experience is chosen with a probability proportional to that experience's perceived *value* in terms of the stored value attributes (see section III-B above). The

[1] Vision sensors here are built by subdividing the visual field into regions and taking average colour values over each region at each timestep. In these experiments a 6x6 grid is used taking the average of the red channel only.

[2] Auditory channels were not used in the examples discussed.

[3] Note that sensor data is not being stored to build the interaction history, only the time-evolving probability distributions from which joint entropy can be estimated are stored.

[4] While this is probably not the most sophisticated model for acting, it is at least tractable.

exact nature of the calculation of *value* is dependent on the nature of the drives and motivations ascribed to the agent. For these experiments we use an internal variable that increases whenever a ball or human face is seen, but decays over time. This is explained in more detail in section *V-C*.

Finally, we introduce a feedback process that evaluates the result of any action taken in terms of whether there was an *increase in value* after the action was executed, and then adjusts the stored value attributes of the candidate experience, from which the action was derived, up or down accordingly. 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 behaviour patterns over time. This can be viewed as ontogenetic development, that is, as a process of change in structure and skills through embodied, structurally coupled interaction [9].

IV. Geometry of Experience

In previous papers [7], [10]–[12] the authors have developed a mathematical geometry of experience that uses Shannon information theory [13] to place experience on a metric space as well as to compare sensorimotor experience using trajectories through projected sensor and motor spaces. The basis is the information metric [14], a measure of the "distance", in terms of information, between two random variables. We use the measure to compare sensorimotor experience over time and across modalities and the following is a brief overview of the relevant aspects.

A. Information Distance

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. For any pair of 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.

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

where $p(x,y)$ is given by the joint distribution of $\mathcal{X}$ and $\mathcal{Y}$.[5]

The *information distance*[6] between $\mathcal{X}$ and $\mathcal{Y}$ is then given by

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

[5] 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.

[6] This satisfies the mathematical axioms for a *metric*:
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).

B. Time-Horizon

Consider any sensor variable $\mathcal{X}$, beginning from a particular moment in time t_0 until a later moment t_0+h ($h>0$), we regard the sequence of values $x(t_0), x(t_0+1), \ldots, x(t_0+h-1)$ taken by an information source $\mathcal{X}$ as 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*

With this definition and that of information distance, we can then compare any sensorimotor variables over the same sized time-horizons, whether from the same sensor at different times, different sensors at the same time or, indeed, different sensors at different times.

C. Experience Metric

We 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

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

where $E = E(t,h)$ and $E' = E(t',h)$ are two experiences of an agent and d is the information distance (see [7], [10]).

V. Experiments

We describe two experiments that explore the possibilities of the model of interaction history discussed. The first evaluates the veracity of the experience space by examining the ability of the model to predict future states of the world with reference only to the metric space of experience. The second shows early steps in using the model to play an interaction game, "peekaboo", with a human partner.

A. Experiment 1: History-based Prediction

Given a robot[7] acting in an environment, how well can it predict future events based on its recent history of experience?

In this experiment the architecture was simplified, removing the developmental feedback loop, to examine the efficacy of using the metric space of experience to locate similar experiences. Two conditions were examined: in the first, the head stayed still while the ball was moved, and in the second a reactive process allowed the head to follow the ball.[8]

The *position* of the ball at the end of each experience was stored with the experience as a value attribute, and the *predicted future position* of the ball was given by the

[7] See Fig. 2. The robot used in this and all other experiments is the Sony AIBO ERS-7. Robot control programming was achieved using URBI-Universal Real-time Behaviour Interface [15].

[8] Simple colour based visual processing allowed the position of a pink ball in the visual field to be located as an (X,Y) position, and the head would reactively move to centre that position.

Fig. 2. Sony Aibo ERS-7, Left: with pink ball, Right: hiding head while playing ”peekaboo”. The camera vision is partially obscured by the arm.

attributes stored with the experiences following the candidate (most similar previous) experience.

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 tagged 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.

B. Experiment 1: Results and Discussion

In Fig. 3, we show a sequence of images from one trial from the first condition where the robot was passive while the ball was moved. The sequence lasts just over 4 seconds and consists of approximately 40 timesteps (1 timestep$\sim$100ms) and 8 experiences (experience granularity[9] of 5) [10].

In the sequence shown and others, the robot required very few examples of a sequence (usually one) before the appropriate 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. 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.

As an illustration of the problem, consider the eighth image in Fig. 3, here the predicted path from 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

[9]Experience granularity denotes the number of timesteps between end-points of successive experiences. A granularity of 1 would store an experience of *horizon* timesteps at every timestep.

[10]Images are saved asynchronously at a rate of approx. 4 per second.. There were approximately 73 experiences at a granularity of 5 timesteps between experiences (about 38 seconds of activity) before the ones shown. Before the images shown, the ball was moved from left to right 4 times and in a clockwise circle once.

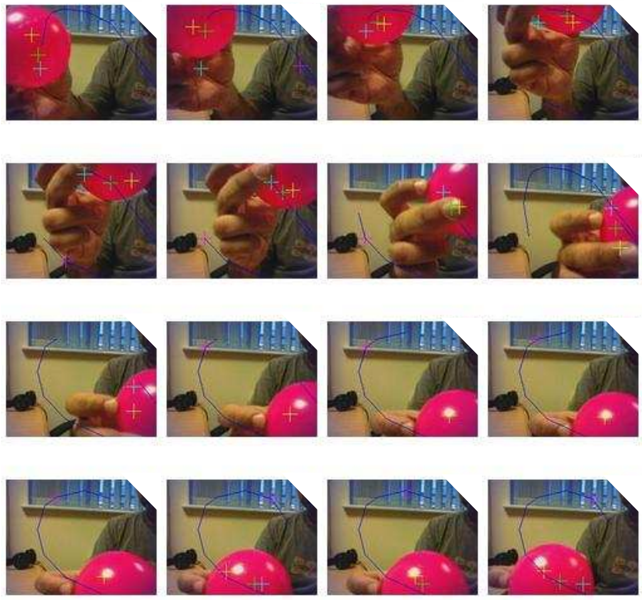

Fig. 3. Series of consecutive images from the Aibo camera showing ball path prediction using a sensorimotor interaction history. The robot does not move its head in this sequence. Images are sequential left to right and top to bottom and 147 images (73 experiences horizon length 40) precede these. The line shows the path prediction for 10 timesteps 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=5 timesteps. Images captured approximately once every 2-3 timesteps.

a cycle back in time that may have been more appropriate, and the fact that the two possible experiences correspond to motions of the ball from opposite sides of a circle contributes to their being “recoding equivalents”[11], only differing in phase. Clearly, one solution to the issue is to provide the mechanism with more information, for instance from proprioception, with which to distinguish experience. The experiment is artificially hampered due there being no motor, active, component to the interaction.

Fig. 4 shows a series of images showing the path prediction in the second condition, where the robot was actively following at the ball with its head. The ball path is generally a small loop starting at and finishing near the centre of the image. This is to be expected as, since we are plotting just the position of the ball *within the image* then, this cannot describe the absolute position of the ball in space. To better assess the result of the experiment, we would need to have the predicted position of the head rather than the ball. Further work, will look into the predictive capabilities of the method with regard to the robot acting as a whole.

C. Experiment 2: Sensorimotor contingencies in an interaction game - Peekaboo

The purpose of this experiment was to investigate whether the development of an enactive interaction history in a

[11]That is, are a small information distance apart.

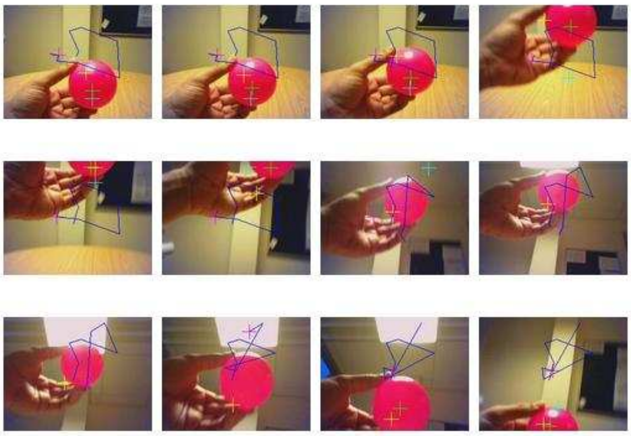

Fig. 4. Series of consecutive images (left to right, top to bottom) from the Aibo camera showing ball path prediction using a sensorimotor interaction history. The robot's head reactively follows the ball in this sequence. Images are sequential left to right and top to bottom and the sequence is approximately 3 seconds long. (See text and Fig. 3 for further notes). Horizon=20, Number of Bins=2, Experience granularity=4 timesteps.

robot could be used for the robot to act appropriately in an interaction that required following a spatio-temporally structured set of "rules", that when followed result in high motivational value. The full architecture was used, with the action feedback loop modifying potential future interaction.

The simple interaction game of *peekaboo* played between adults and babies or young children was taken as a model. The game consists of a repeated cycle of an initial contact, disappearance, reappearance, and acknowledgment of renewed contact [16]. Bruner and Sherwood suggest that the peekaboo game may provide scaffolding for further interaction and learning [16] and as such is useful in studying the development of interaction capabilities in a robot in a social environment.

Bruner and Sherwood also suggest that the peekaboo game itself may emerge from the exploitation of innate tendencies or motivations in the child and we model important aspects of potential precursors to this game as actions, drives and motivations of the cognitive model of the robot. Specifically, the robot gains "pleasure" (increase in internal variable 1) in seeing a face, however if the face is lost, it has a rising "expectation" (internal variable 2) of seeing the face again, and the "pleasure" in seeing the face at a later time is increased by the value of that expectation. The atomic actions implemented for the selection mechanism were: 1) move head up, 2) down, 3) left, 4) right and 5) hide/reveal head.
The robot is preprogrammed with abilities to recognize a generalized face[12] and this yields sensory variables indicating the position of the face in the visual field.

[12]Implemented using Intel OpenCV HAAR Cascades [17].

D. Experiment 2: Results and Discussion

Thus far we have completed a basic feasibility study with one of the authors interacting with the robot. The results tend to show that the robot, after a period of random movement does start to engage in repeated cycles of behaviour, Fig. 5. If the robot were not to hide its face, it would have long periods of seeing the face which do not result in high motivational value (internal variable 1), instead the robot generates intermittency in seeing the face by hiding its own face resulting in high motivational value when the face is next seen. This often includes cycles of hiding and revealing the face, as shown in Fig. 5.

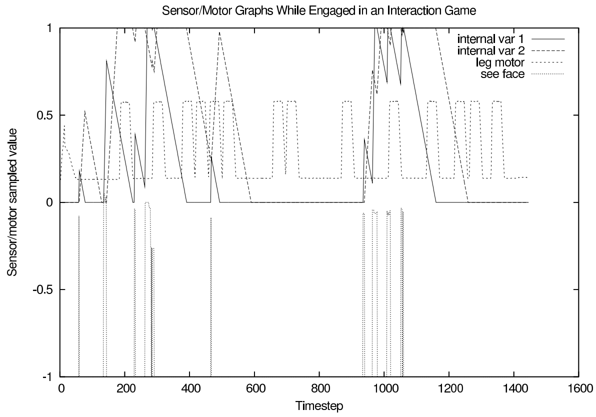

Fig. 5. Time series of motor and sensor values showing engagement of robot in peekaboo game. The bottom part of the graph shows when the face is seen and the two internal variables are shown varying in response to this. The peaks in the leg motor trace indicate when the robot is hiding its head with its foreleg.

Fig. 6 shows the value assigned to experiences and how these change over time. It is clear that only a few experiences are regularly selected and thus modified over time, increasing and decreasing in value. The final metric space of experiences is depicted in Fig. 7, and indicates that the experience space has a consistent (non-random) structure with definite peaks that correspond to those few experiences that become present candidate experiences.

VI. Discussion and Future Directions

The positive results from the experiments using the experience space to predict future experiences indicate that the method of information distance has the potential of forming the basis of an interaction history, particularly if the whole embodied experience of the robot is taken into account. However, mechanisms may be needed to disambiguate the experience in the space when there are many experiences to select from. Steps towards this are made in using "value" to test candidate experiences against each other, however, other mechanisms might be considered, e.g. finding exemplary experiences by grouping near experiences. Further, it is also

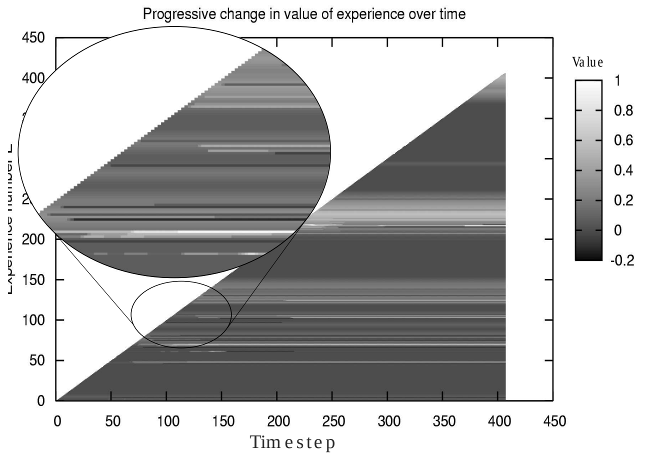

Fig. 6. Graph showing the "value" (as shade) assigned to experiences (on vertical axis), and how this progresses over time as values are changed while the robot actively reconstructs its experience space. The zoomed in region shows individual experiences changing in value. Note that the triangular shape is due to new experiences being added over time, and that most experiences do not change in their values.

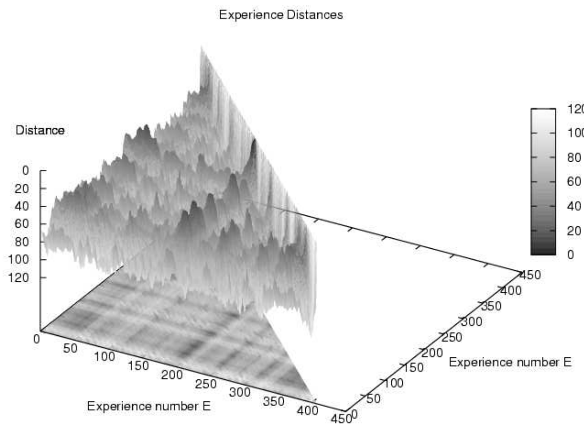

Fig. 7. Depiction of experience space at the end of the run shown in Fig. 6. The axes in the plane are the experiences being compared while the height indicated their experiential information distance. The black peaks are low information distances and indicate "similar" experiences. It is these experiences that provide candidates from which to select action.

clear that experiences of different time horizon sizes will be needed to anticipate experience on different timescales.

There are also good indications that the method of choosing action based on "value" can be useful in choosing between many potentially similar experiences, however, how this value is assigned and modified will need to be made more sophisticated in order to better assign credit to appropriate action, and to handle multiple, potentially conflicting, goals. Similar comments apply to the simple method of action selection. In a more complex environment an action selection mechanism that can deal with appropriate action in a particular context, and that can deal with parallel and temporally extended actions and behaviours would be needed.

Future directions for the research will include investigating more generic approaches to ascribing motivation to artificial agents in order to select experience and action, for example sensorimotor contingencies, drives for comfort (predictability of environment) or novelty (see for example [18]). Further work will be conducted on the peekaboo game as a testbed in which to study the development of interactive behaviour.

REFERENCES

[1] H. R. Maturana and F. J. Varela, *The Tree of Knowledge: The Biological Roots of Human Understanding*. Boston: New Science Library (Shambhala), 1987.

[2] W. J. Clancey, *Situated Cognition: On human knowledge and computer representations*. Learning in doing: Cognitive and compuational perspectives, Cambridge University Press, 1997.

[3] K. Dautenhahn and T. Christaller, "Remembering, rehearsal and empathy - towards a social and embodied cognitive psychology for artifacts," in *Two Sciences of the Mind: Readings in cognitive science and consciousness* (S. O'Nuallain and P. M. Kevitt, eds.), pp. 257–282, John Benjamins North America Inc., 1996.

[4] E. Tulving, *Elements of episodic memory*. Oxford: Clarendon Press, 1983.

[5] C. L. Nehaniv, K. Dautenhahn, and M. J. Loomes, "Constructive biology and approaches to temporal grounding in post-reactive robotics," in *Sensor Fusion and Decentralized Control in Robotics Systems II (September 19-20, 1999, Boston, Massachusetts), Proceedings of SPIE* (G. T. McKee and P. Schenker, eds.), vol. 3839, pp. 156–167, 1999.

[6] E. T. Rolls, *The Brain and Emotion*. Oxford University Press, 1999.

[7] C. L. Nehaniv, "Sensorimotor experience and its metrics," in *Proc. of 2005 IEEE Congress on Evolutionary Computation*, 2005.

[8] A. Meltzoff and M. Moore, "Explaining facial imitation: a theoretical model," *Early Development and Parenting*, vol. 6, pp. 179–192, 1997.

[9] M. Lungarella, G. Metta, R. Pfeifer, and G. Sandini, "Developmental robotics: A survey," *Connection Science*, vol. 15, no. 4, pp. 151–190, 2004.

[10] C. L. Nehaniv, N. A. Mirza, K. Dautenhahn, and R. te Boekhorst, "Extending the temporal horizon of autonomous robots," in *Proc. of the 3rd International Symposium on Autonomous Minirobots for Research and Edutainment (AMiRE2005)* (K. Murase, K. Sekiyama, N. Kubota, T. Naniwa, and J. Sitte, eds.), pp. 389–395, Springer, 2006.

[11] N. A. Mirza, C. Nehaniv, R. te Boekhorst, and K. Dautenhahn, "Robot self-characterisation of experience using trajectories in sensory-motor phase space," in *Proc. of Fifth International Workshop on Epigenetic Robotics: Modeling Cognitive Development in Robotic Systems (EpiRob2005)*, pp. 143–144, Lund University Cognitive Studies, 2005.

[12] N. A. Mirza, C. L. Nehaniv, K. Dautenhahn, and R. te Boekhorst, "Using sensory-motor phase-plots to characterise robot-environment interactions," in *Proc. of 6th IEEE International Symposium on Computational Intelligence in Robotics and Automation (CIRA2005)*, pp. 581–586, 2005.

[13] C. E. Shannon, "A mathematical theory of communication," *Bell Systems Technical Journal*, vol. 27, pp. 379–423 and 623–656, 1948.

[14] J. Crutchfield, "Information and its metric," in *Nonlinear Structures in Physical Systems - Pattern Formation, Chaos and Waves* (L. Lam and H. Morris, eds.), pp. 119–130, New York: Springer-Verlag, 1990.

[15] J.-C. Baillie, "Urbi: Towards a universal robotic low-level programming language," in *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems 2005*, 2005. http://www.urbiforge.com.

[16] J. S. Bruner and V. Sherwood, "Peekaboo and the learning of rule structures," in *Play: Its Role in Development and Evolution* (J. Bruner, A. Jolly, and K. Sylva, eds.), pp. 277–285, New York: Penguin, 1975.

[17] OpenCV, "Open computer vision library." http://sourceforge.net/projects/opencvlibrary/ (GPL), 2000.

[18] P.-Y. Oudeyer, F. Kaplan, V. V. Hafner, and A. Whyte, "The playground experiment: Task-independent development of a curious robot," in *Proc. of the AAAI Spring Symposium on Developmental Robotics, 2005* (D. Bank and L. Meeden, eds.), pp. 42–47, 2005.

B.1.2 Peekaboo: Effect of experience length on the interaction history driven ontogeny of a robot.

N. A. Mirza, C. L. Nehaniv, K. Dautenhahn, and R. te Boekhorst. Peekaboo: Effect of experience length on the interaction history driven ontogeny of a robot. In *Proceedings the of 6th International Conference on Epigenetic Robotics*, pages 71–78, Paris, France, 20-22 September 2006. Lund University Cognitive Studies. 2006.

Peekaboo: Effect of Experience Length on the Interaction History Driven Ontogeny of a Robot

Naeem Assif Mirza **Chrystopher Nehaniv** **René te Boekhorst** **Kerstin Dautenhahn**

Adaptive Systems Research Group, School of Computer Science, University of Hertfordshire, College Lane, Hatfield, Hertfordshire. United Kingdom. AL10 9AB

{N.A.Mirza,C.L.Nehaniv,R.teBoekhorst,K.Dautenhahn}@herts.ac.uk

Abstract

The game peekaboo, ordinarily played between an adult and baby, is used as a situation where a robot may develop social interaction skills such as rhythm, timing and turn taking, using its experience and history of interactions over different temporal horizons. We present experiments using a robot that explore the length of experiences in an architecture that selects action based on a metric space consisting of previous experience and feedback from the environment. Results show that sequences of interactions that allow the robot to play the game successfully emerge from the interplay between environmental or social feedback and experience of various lengths.

1. Introduction

One of the main challenges faced in building agents embedded in a social environment is how they can make use of their experience and history of interaction to modulate future action in a meaningful way and to be further shaped by that action. We take the view that appropriate mechanisms, while based in innate abilities, should largely develop through ontogeny. Our approach is to conduct experiments on a physical robot (see figure 1) to examine these mechanisms for development.

In the study of the ontogeny of social interaction and turn-taking in artificial agents, it is instructive to look at the kinds of interactions 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" and in general consists of a repeated cycle of an initial contact[1], disappearance, reappearance, and acknowledgement of renewed contact (Bruner and Sherwood, 1975). Bruner and Sherwood note that, while the peekaboo game itself

Figure 1: *Aibo playing "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.)

emerges from the exploitation of an innate tendency in the child that is rewarded by pleasure in responsiveness, the game is highly rule bound and needs to be learnt.

Peekaboo is a common game played by very young children[2] with adults. The contingent, temporal structure of the game makes it useful as a tool to better understand the role of interaction as a possible mechanism to ground robot ontogeny in human-robot interaction. The child must develop some anticipation of what might happen in the future, and, moreover, the meeting of this expectation (or indeed, failure to meet) is where the fun and interest inherent in the game comes from.

The rhythm and timing of the interaction are crucial and, Bruner and Sherwood suggest 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, pp424-425).

The temporal structure of the peekaboo game suggests that a robot control or cognitive architecture needs to take into account the history of interaction. We describe an architecture where an embodied robotic agent can make use of an interaction history to guide ontological development to act appro-

[1]Initial contact is usually face-to-face mutual looking (Bruner and Sherwood, 1975).

[2](Bruner and Sherwood, 1975) studied 7 month old to 17 month old children but note that the game is played by younger children still.

priately in a changing environment. The direct sensorimotor motor history of the agent is used to create grounded experiences of different lengths which can be compared with one another using a metric measure based on the information distance between them. The agent acts on the basis of its experiences and the choice of action depends, in part, on the feedback of reward from the environment. This architecture was initially explored in (Mirza et al., 2006), where the efficacy of the experience space was verified by using the history to predict the future position of a ball.

To relate experiences with other experiences in an interaction history, we use information distance measures (Shannon, 1948, Crutchfield, 1990) and a mathematical concept of experience and the relations between them. These are defined in (Nehaniv, 2005) and reviewed in Section 2.4. Information distance related techniques have been successfully used in the past, for instance, to compare behaviours from the perspective of the agent (Mirza et al., 2005, Kaplan and Hafner, 2005) and for an agent to infer a model of its own sensory and actuator apparatus by acting in the environment (Olsson et al., 2005). This suggests that behaviour can be guided by moving in a continually constructed space of experiences by selecting appropriate actions that will move the agent closer to desired experiences.

We emphasise an ontogenetic developmental approach (Lungarella et al., 2004, Blank et al., 2005) to acquiring appropriate behaviour, in that, the structures controlling action are modified by interaction and experience and new skills are acquired. A new feature of our approach is the growth and exploitation of the developing agent's (metric) space of experiences driving its ontogeny in interaction with its environment.

This paper continues by describing in further detail the model of interaction history, the metric space of experience and implementation in a physical robot. We then describe experiments where we investigate the effect of temporal scale (horizon) of experience on the ability of the robot to develop in playing the game. We conclude the paper with the results of the experiments and a discussion of the strengths and limitations of the current model, and outline how future research can further improve the models discussed.

2. Model of Interaction History

In developing a model of interaction history we start out by considering what such a history might be, and present a working definition. We then describe the model in outline and go on to explain its key parts, namely: the metric space of experience, the action selection mechanism and the motivational subsystem.

2.1 Interaction History

We use a working definition of an *interaction history* as:

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

The key aspects of this definition are:

- *Temporal extension*: experiences are associated to episodes of particular duration in terms of events experienced by the agent. The horizon[3] of an agent extends into the past (including all previous experience available to the agent) and also into the future in terms of prediction, anticipation and expectation.
- *Dynamic 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 representational (i.e. recorded in terms of imposed representations) and is grounded in the sensorimotor experience of the agent.
- *Remembering, manifest as action*: "memory" consists not of static representations of the past that can be recalled with perfect clarity, but rather is the result of an accumulation of interaction with the environment and this history of interaction is revealed as current and future action. See for example (Rosenfield, 1988, Dautenhahn and Christaller, 1996).

2.2 An Interactive History Architecture

We describe a computational model (Figure 2) that demonstrates how such interaction histories can be explicitly integrated into the control of a robot. The basic architecture consists of processes to acquire sensory and motor data from the robot as it acts in the environment (see Section 2.3), from this a metric space consisting of past interaction experiences is constructed (see Section 2.4). A process then selects past experiences near (i.e. with low information distance) to the current experience (see Section 2.5). The selection is also based on the values of internal variables that change according to a motivational system (see Section 2.6). The action following the chosen past experience becomes the next action of the agent. Finally, there is an internal feedback process that adjusts the values of internal variables associated with any experience when it has been used

[3] Horizon has a different technical meaning when we talk of the *horizon length of an experience* as detailed in Section 2.4

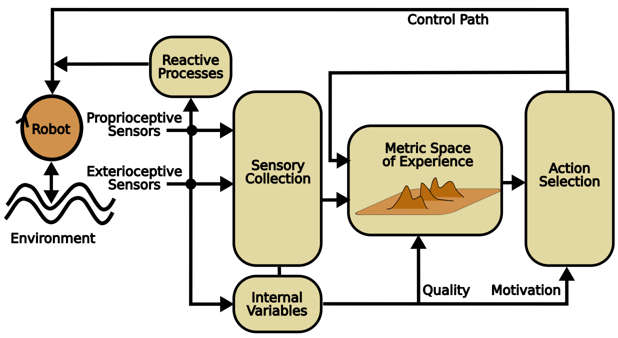

Figure 2: Interaction history based control architecture.

to select future action, making it more or less likely to be chosen in the future.

There are many potential architectures that take history of action and interaction into account, including top-down deliberative architectures such as Soar (Nuxoll and Laird, 2004), connectionist systems that have memory, for instance Elman networks or recurrent neural networks (Rylatt and Czarnecki, 2000) and certain behaviour oriented control systems combined with learning (Matarić, 1992, Michaud and Matarić, 1998). Our model is not deliberative as no overall plan is constructed, and it makes history explicit and inspectable unlike neural network approaches in general. Most behaviour 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 behaviour being selected (or indeed, the action being selected), but in terms of the sensorimotor history.

2.3 Sensory and Internal Variables

The sensory information available to the robot[4] 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). The actual variables used in this implementation are summarised in (Table 1), with further discussion of internal variables in Section 2.6 and Appendix A.

All the variables are treated as "random variables" with local stationarity, for which we can estimate the probability distributions and entropy for the purpose of calculating information distance and the experience metric. See Section 2.4. We also use certain of these variables to indicate "quality" and in these cases, the instantaneous values of those variables at

[4]Sampling is done at regular intervals (between 100-120ms in the experiments here). Vision sensors are built by subdividing the visual field into regions and taking average colour values over each region at each timestep. In these experiments a 3x3 grid over the image is used taking the average of the red channel only.

Table 1: Sensors and Internal Variables

Type	Examples	Total
Exterioceptive	IR-distance, Buttons	15
Proprioceptive	Joint positions,	18
Visual	Average colour values in a 3x3 grid over image	9
Internal	Face position, ball position, desire to see a face	10

the end time point of the experience is attached to the experience.

2.4 Experience Space

The *metric space of experience* is constructed from "experiences" of a particular horizon length (in timesteps) with relative positions in the space determined by the information distance between them.

We formalise 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 $\mathcal{X}^n_{t,h}$ is the random variable estimate from the sequence of values taken by a sensor n from time t to $t+h$ taken from the set of all sensorimotor inputs available to the agent. A metric on experiences of temporal horizon h is then defined as

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

where $E = E(t,h)$ and $E' = E(t',h)$ are two experiences of an agent and d is the information distance between two random variables $\mathcal{X}$ and $\mathcal{Y}$ given by $d(\mathcal{X},\mathcal{Y}) = H(\mathcal{X}|\mathcal{Y}) + H(\mathcal{Y}|\mathcal{X})$. The information distance satisfies the axioms of a metric and can be estimated from the probability distributions[5] of the sampled, discretised variables. See (Nehaniv, 2005) for proofs and discussion.

2.5 Action Selection and Development

While an experience space can be built without much difficulty, the challenge is how to have experience modulate future action in a meaningful way and to be further shaped by that action. To achieve this goal, a simple mechanism is adopted whereby the robot can execute one of a number of "atomic" actions (or no action) at every timestep (see table 2). Each action takes 2 seconds or less to execute and the re-centre head action is duplicated to offset the two actions which take the head away from the centre. A record of actions executed by the robot at any time is kept to facilitate the action-selection based on history of experience.

[5]Note that the discretised (binned) values of all variables at all time intervals are stored in order to be able to estimate the joint distribution with other (new) experiences.

Table 2: Actions

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

To choose an action based on experience, a number of *candidate experiences* from the experience space near to (that is with short information distance to) the current experience are selected, and one chosen according to:

$$p_{E^n} \propto \frac{Q_{E^n}}{D(E^n, E^{current})} - C \qquad (1)$$

where p_{E^n} is the probability of choosing a candidate past experience E^n with quality Q_{E^n}, taken from the set of K experiences $\{E^1, \ldots, E^K\}$ in the neighbourhood of the current experience $E^{current}$. The exact nature of the calculation of *quality* is dependent on the nature of the drives and motivations ascribed to the agent (see section *2.6* and Appendix *A*).

The next action that was executed following the chosen past experience is then the action to be executed next.

If none of the candidate experiences is chosen, then a random action is executed. This has the advantage of emulating body-babbling, i.e. apparently random body movements that have the (hypothesised) purpose of learning the capabilities of the body in an environment (Meltzoff and Moore, 1997). Early in development, there are fewer 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. Additionally, with a small probability reflected by the constant C above, the robot may still choose a random action as this may potentially help to discover new, more salient experiences.

Finally, we introduce a feedback process that evaluates the result of any action taken in terms of whether there was an *increase in quality* after the action was executed, and then adjusts the quality of the candidate experience, from which the action was derived, up or down accordingly. 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 behaviour patterns over time. This can be viewed as ontogenetic development, that is as a process of change in structure and skills through embodied, structurally coupled interaction (Lungarella et al., 2004).

Our approach uses temporally extended experience rather than instantaneous state[6]. 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 (this is particularly important when there is an interaction partner) as is the case with most reinforcement learning paradigms (Sutton and Barto, 1998) and in particular with approaches that do not use a model, for example Q-learning.

Related work in the multi-agent domain (Arai et al., 2000) has agents in a grid world acquiring coordination strategies, and uses a fixed-length episodic history expressly to counter the MDP assumption. However, that 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, nor locate them in a metric space.

2.6 Environmental feedback

We make use of feedback from the environment as actions are executed, and define certain internal variables and their dynamics such that they provide feedback appropriate for the peekaboo game (noting that an appropriate temporal arrangement of actions is still necessary to actually play the game). This can be seen as building in innate drives and motivations in the robot that underly and scaffold the learning of the rules of interaction games, in a way analogous to inherited drives and motivations in human babies.

To provide appropriate feedback, we require a high value for motivation when a face is seen following a period where there has been no face seen. Three internal variables are used to model this: f indicating when a face is seen, m the motivational value that is used as the *quality* of experience, and d the desire to see a face when one is not seen. The exact nature of the dynamics is determined by 6 parameters encoding rates of decay, increase and feedback of f and d. For details see Appendix *A*.

3. Experiments

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. Secondly, the hypothesis that the horizon length of experience would affect the ability to learn 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

[6] that is the instantaneous values of the sensory variables

make useful comparisons to the history. If it is too long, then the interesting part of the interaction becomes lost in the larger experience.

3.1 Implementation and Setup

The architecture was implemented using using URBI (Baillie, 2005) and Java on a Sony Aibo ERS-7 robot dog and a desktop computer. The system runs online with telemetry data being sent over wireless to a desktop approximately every 120ms where the metric space of experience is constructed and used in action selection.

The robot stays in a "sitting" position throughout the experiments with the forelegs are free to move, facing a picture of a face (see Figure 1) at a fixed distance of 40cm. 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. Early experiments where the robot faced a human interaction partner are presented in (Mirza et al., 2006) and this is also the subject of future experiments.

For the purposes of these trials, we define peekaboo-like behaviour to have occurred when face detection has been lost and then regained (one or more times) resulting in a maximum value for the motivational variable m. The duration of the sequence being taken from the point of the first loss of face through to the last point at which high motivation can be sustained without a break in the sequence.

We ran 6 trials of 2 minute duration for each horizon length of 8, 16, 32, 64 and 128 timesteps (0.96, 1.92, 3.84, 7.68 and 15.36 seconds respectively). For comparison, a further 6 trials were run where the action selection was random and not based on history. In each of the trials the metric space started unpopulated.

4. Results

Table 3 summarises the results of 36 trial runs, while Figure 3 shows, for selected trials, time-series graphs of the motivational variables coupled with the actions taken. Peekaboo behaviour, involving hiding the head, was seen in 18 of the 36 runs. All of the trials using random action selection showed some peekaboo behaviour, although it was intermittent and not regular (see figure 3A for example). All but one of the horizon size 8 trials, and all but two of horizon size 16, also showed peekaboo, however, there were longer periods of repeated behaviour. Figure 3A (horizon size 8) shows the best example of an extended period of peekaboo behaviour; the repeat period is approximately 42 timesteps or 5 seconds, and the episode continues for around 640 timesteps (76 seconds). During this episode the head is hidden to the left and right and this is interspersed with head-centring actions. Through all of these episodes periods of no action serve to alter the timing of the cyclic periods.

Of the longer horizon length (32, 64 and 128) trials, three showed peekaboo behaviour, but three also showed an emergent behaviour which resulted in high motivation, see Figure 3C for an example. Here the robot stares ahead at the face while intermittently waving. Due to the way that the robot was sat during some of these trials the robot was shaken slightly as the front arm finished the wave and rested on the hind leg, causing a momentary loss of face detection. Given the sensitivity of the motivational system, this was enough, when repeated, for the dynamics to result in increased desire d and therefore high motivation m.

5. Discussion and Future Work

All of the trial runs where only random actions were selected resulted in some episodes of high motivational value (m). It is likely that this is due to a very sensitive motivational system[7] combined with a range of actions, most of which would result in some loss of face detection. However, to see longer periods of high motivation, some controlled behaviour must be selected (as a contrary example see Figure 3F where no peekaboo-like dynamics are seen). Cyclic behaviour with the long peekaboo-like sequences of repeated action is only seen in the experience-driven trials.

In some of the experience-driven trials repeated behaviour was seen that could have resulted in high motivation were the head pointed forward, however, a single action turned the head away, and experience alone was not able to re-centre the head. On one occasion however, when the head was re-centred (randomly) then the experience space allowed a resumption of the peekaboo sequence (see figure 3E). It is possible that if each trial had a longer duration, then the experience space would be richer and recentring behaviour would be selected. This also may point to a reason why the trials using longer horizon lengths performed poorly: appreciation of current state may be necessary to notice that the head is not pointing forward (for instance) and this may be easier with a shorter time horizon.

The best of the cyclic behaviour was seen in the experience-driven trials of horizon size 8 and 16 timesteps (approx 1 and 2 seconds respectively). This result indicates that it is necessary to have a

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

Table 3: *Experiment Summary.* Duration and period in timesteps (ts) of peekaboo (pkb) behaviour for each trial. Also noted is where high m is attained with an alternative, emergent sequence.

Run	**Random** length/period	**Horizon 8** length/period	**Horizon 16** length/period	**Horizon 32** length/period	**Horizon 64** length/period	**Horizon 128** length/period
1	120ts / 40ts	180ts / 45ts	260ts / 40ts	none	Waving pkb 400ts	none
2	220ts / 55ts	150ts / 40ts	none	none	none	none
3	220ts / 45ts	*fig 3A*, 640ts/42ts	140ts/45ts,200ts / 50ts	*fig 3F*, none	none	100ts / 40ts
4	200ts / 60ts	130ts / 45ts 150ts / 70ts	*fig 3E*, 260,240ts/40ts repeated sequence	none	none	none
5	160ts / 50ts	none	Waving emergent pkb 150ts	*fig 3C* Waving pkb 540ts / 47ts	*fig 3D*, 160,100,140ts / 40ts	120ts / 40ts
6	*fig 3B* 80,140ts / 40ts	250ts / 42ts	120ts / 40ts	Waving pkb 840ts / 47ts	none	none

short time-horizon, and this may be related to the length of single actions (about 2 seconds), and thus the natural period[8] of the cyclic behaviour. A reason why this may be the case is that, to bootstrap the initial repetitive behaviour, 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.

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 with parts of longer term episodic experience, and the current short experience being given 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 short horizons 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 kinds 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 determination of which will take into account motivational dynamics, value of experience, and possibly rates of change of experience distances.

These particular experiments do not have any interaction from the partner's side and so are lacking a vital part of the interaction. The motivational dynamics compensate for this by providing a reward landscape based only on internal factors and the single external stimuli of a face. However, we argue that the interaction history can be extended to a fully interactive scenario by, for instance having the interaction partner modulate both the external stimulus (the presentation of the face) as well as, potentially, the reward signal that interacts with the motivational dynamics. Given that the robot's actions can in some way affect the behaviour of the partner (e.g. bark excitedly when an internal variable reached maximum), then the interaction history could be used as part of a full interaction.

The motivational system used is specific to the peekaboo scenario, and while it potentially gives useful insights into motivational dynamics for other scenarios, is not generally applicable. Additionally it is clear that the system is overly sensitive with high motivational value being reached very easily through a wide range of interactions. As an alternative it would be useful to explore the balance between novelty and mastery drives as in, for example (Oudeyer et al., 2005), as the basis of a more general motivation system. Moreover, basing novelty and mastery directly on the structure of the experience space as it develops through interaction would ground these notions in the sensorimotor history of the agent.

Finally, we conclude that the architecture is able to direct future action of an agent based on previous experience and that the horizon length of experience plays an important role in the types of interaction that can be engaged. 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 extremely limited and simplistic and this combined with the short experiment lengths and the over-sensitive motivational system suggests various directions for improvement.

[8] Note that the motivational system itself does not dictate this period as any cyclic behaviour of period up to 19 seconds can result in high values of m.

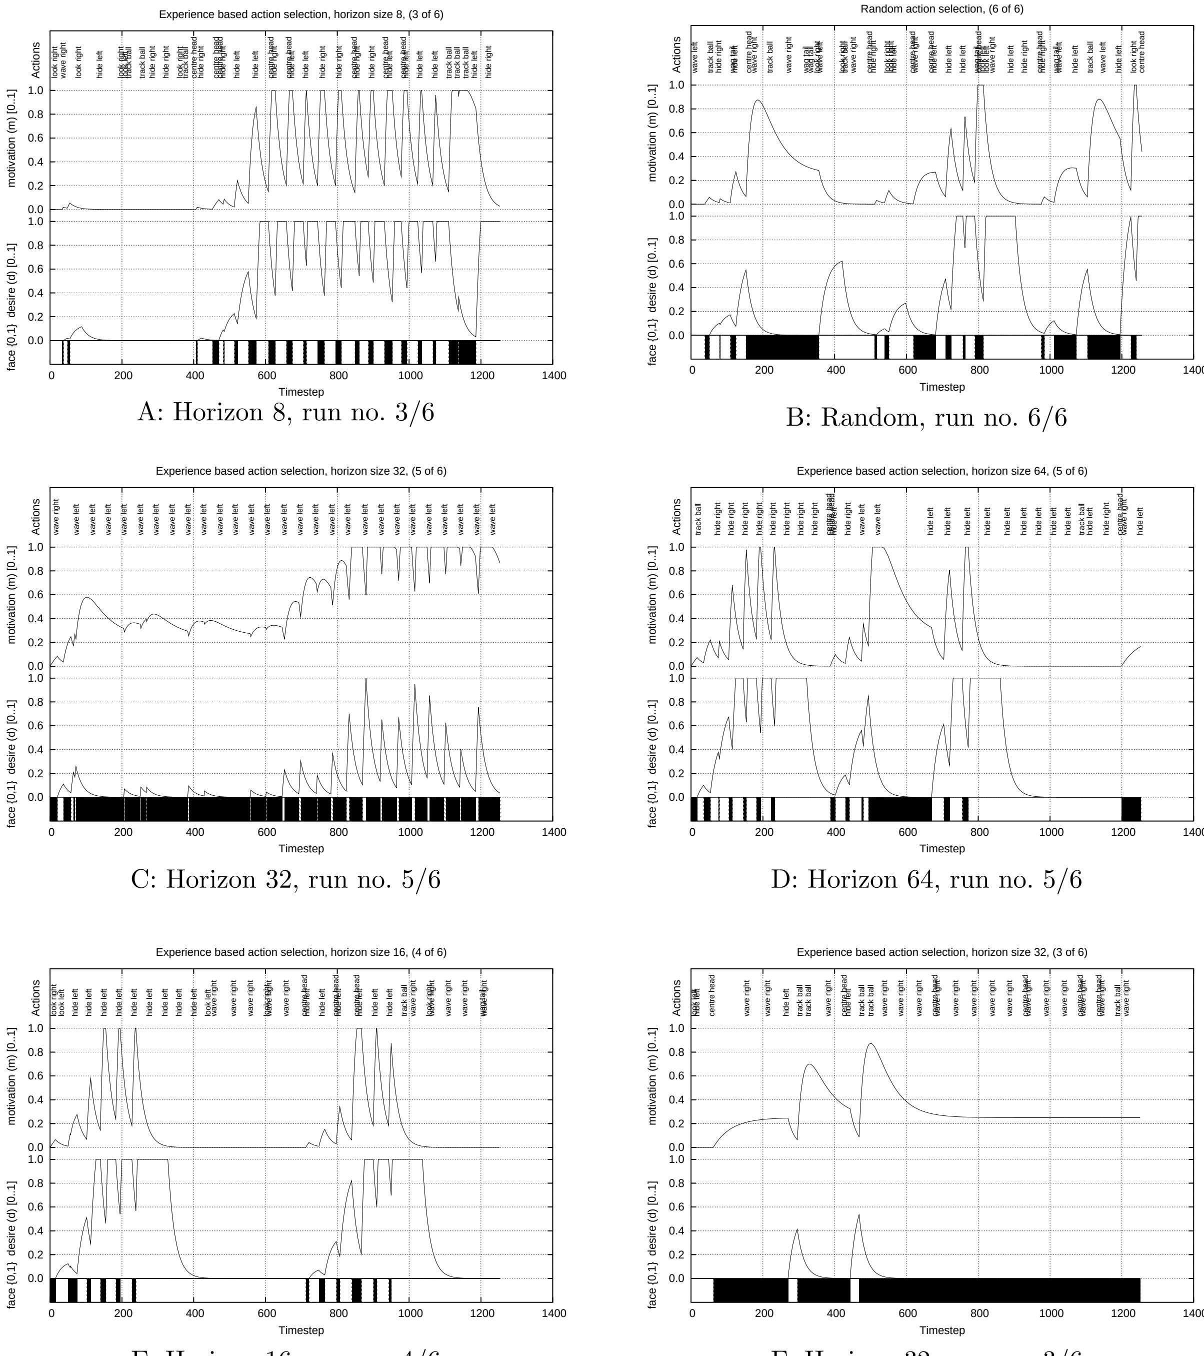

A: Horizon 8, run no. 3/6

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

Figure 3: Motivational dynamics and actions for selected 2 minute 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: 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 get periods of high value. **C:** *Emergent behaviour resulting in high m and d.* Horizon size 32. Dynamics generate high value when 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 high value, this maybe 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 recentred. **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.

Appendix A - Motivational Dynamics

Firstly, the agent possesses a binary meta-sensor f that is a result of processing the visual sensors (image) to locate a generalised human face shape in the image, if one exists[9]. This is smoothed to remove short gaps ($< 50ms$).

Secondly, 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 2.

Finally, the overall motivation m, also constrained in the range [0,1] and increases when $f = 1$ 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 3.

$$\Delta d = \begin{cases} \alpha_1 m - \delta_1(1-m)d & \text{if } f = 0, \\ -\delta_2 d & \text{if } f = 1. \end{cases} \tag{2}$$

$$\Delta m = \begin{cases} -\delta_3 m & \text{if } f = 0, \\ \alpha_2 d + \beta(C_{max} - m) & \text{if } f = 1. \end{cases} \tag{3}$$

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

The parameters of the dynamics equations are shown below along with the values used in the experiments. These values were chosen by trial and error.

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

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 are grateful to Martin Gruendl for permission to use the average female face from the Beautycheck project.

References

Arai, S., Sycara, K., and 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*, pages 125–135.

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

Blank, D., Kumar, D., Meeden, L., and Marshall, J. (2005). Bringing up robot: Fundamental mechanisms for creating a self-motivated, self-organizing architecture. *Cybernetics and Systems*, 36(2).

Bruner, J. S. and Sherwood, V. (1975). Peekaboo and the learning of rule structures. In Bruner, J., Jolly, A., and Sylva, K., (Eds.), *Play: Its Role in Development and Evolution*, pages 277–285. New York: Penguin.

Crutchfield, J. (1990). Information and its metric. In Lam, L. and Morris, H., (Eds.), *Nonlinear Structures in Physical Systems - Pattern Formation, Chaos and Waves*, pages 119–130. Springer-Verlag, New York.

[9]Implemented using Intel OpenCV HAAR Cascades (OpenCV, 2000).

Dautenhahn, K. and Christaller, T. (1996). Remembering, rehearsal and empathy - towards a social and embodied cognitive psychology for artifacts. In Seán Ó'Nualláin, Paul Mc Kevitt, and Eoghan Mac Aogáin, (Eds.), *Two Sciences of the Mind: Readings in cognitive science and consciousness*, pages 257–282. John Benjamins North America Inc.

Kaplan, F. and Hafner, 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)*, pages 129–134. IEEE.

Lungarella, M., Metta, G., Pfeifer, R., and Sandini, G. (2004). Developmental robotics: A survey. *Connection Science*, 15(4):151–190.

Matarić, M. J. (1992). Integration of representation into goal-driven behaviour-based robots. *IEEE Transactions on Robotics and Automation*, 8(3):304–312.

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

Michaud, F. and 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., and te Boekhorst, R. (2005). Using sensory-motor phase-plots to characterise robot-environment interactions. In *Proc. of 6th IEEE International Symposium on Computational Intelligence in Robotics and Automation (CIRA2005)*, pages 581–586.

Mirza, N. A., Nehaniv, C. L., Dautenhahn, K., and te Boekhorst, R. (2006). Interaction histories: From experience to action and back again. In *Proceedings the of 5th International Conference on Development and Learning (ICDL 2006)*. ISBN 0-9786456-0-X.

Nehaniv, C. L. (2005). Sensorimotor experience and its metrics. In *Proc. of 2005 IEEE Congress on Evolutionary Computation*.

Nuxoll, A. and Laird, J. E. (2004). A cognitive model of episodic memory integrated with a general cognitive architecture. In Lovett, M., Schunn, C., Lebiere, C., and Munro, P., (Eds.), *Proceedings of the Sixth International Conference on Cognitive Modeling*, pages 220–225. Lawrence Erlbaum Associates.

Olsson, L., Nehaniv, C. L., and Polani, D. (2005). From unknown sensors and actuators to visually guided movement. In *Proceedings the of International Conference on Development and Learning (ICDL 2005)*, pages 1–6. IEEE Computer Society Press.

OpenCV (2000). Open computer vision library. http://sourceforge.net/projects/opencvlibrary/ (GPL).

Oudeyer, P.-Y., Kaplan, F., Hafner, V. V., and Whyte, A. (2005). The playground experiment: Task-independent development of a curious robot. In Bank, D. and Meeden, L., (Eds.), *Proc. of the AAAI Spring Symposium on Developmental Robotics, 2005*, pages 42–47.

Pea, R. D. (2004). The social and technological dimensions of scaffolding and related theoretical concepts for learning, education and human activity. *The Journal of the Learning Sciences*, 13(3):423–451.

Rosenfield, I. (1988). *The Invention of Memory: A New View of the Brain*. Basic Books: New York.

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

Shannon, C. E. (1948). A mathematical theory of communication. *Bell Systems Technical Journal*, 27:379–423 and 623–656.

Sutton, R. S. and Barto, A. G. (1998). *Reinforcement Learning: An Introduction*. MIT Press.

B.2 Interactive Humanoid

B.2.1 Perception of robot smiles and dimensions for human-robot interaction design

Mike Blow, Kerstin Dautenhahn, Andrew Appleby, Chrystopher L. Nehaniv, and David Lee. In *The 15th IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN 06)*, pages 469–474, Hatfield, UK, 6-8 September 2006.

Perception of Robot Smiles and Dimensions for Human-Robot Interaction Design

Mike Blow, Kerstin Dautenhahn, Andrew Appleby, Chrystopher L. Nehaniv, David C. Lee

***Abstract*— As robots enter everyday life and start to interact with ordinary people the question of their appearance becomes increasingly important. Our perception of a robot can be strongly influenced by its facial appearance. Synthesizing relevant ideas from narrative art design, the psychology of face recognition, and recent HRI studies into robot faces, we discuss effects of the uncanny valley and the use of *iconicity* and its relationship to the self-other perceptive divide, as well as *abstractness* and *realism*, classifying existing designs along these dimensions. A new expressive HRI research robot called KASPAR is introduced and the results of a preliminary study on human perceptions of robot expressions are discussed.**

I. Motivations

It is an exciting time in robotics. Personal service robots, so long the science fiction dream, are becoming reality and are for sale to general consumers. Currently their uses are limited, but capabilities are improving, costs are coming down and sales are growing. In addition robots are finding a new place in society as toys, artificial pets [20], security guards, teachers [10], tour guides [24] and in search and rescue. They are finding use in areas as diverse as autism therapy [22], space exploration and research into cognition and biological systems [23].

A. RobotCub

One such research project that we are involved in at Hertfordshire is RobotCub, a multinational European project to build a humanoid child-size robot for use in embodied cognitive development research [23]. The RobotCub consortium consists of 11 core partners from Europe with collaborators in America and Japan, and the institutions involved are each working on specific areas of the robot design, engineering, developmental psychology and human-robot interaction. The software APIs and hardware plans will be published under open-source licenses, with the aim of creating a community using a common platform for robotic and cognitive research.

B. Designing Robots for Users

A previous study of people's expectations of a robot companion indicated that a large proportion of the participants in the test were in favour of a robot companion, especially one that could communicate in a human-like way [6]. Human-like behaviour and appearance were also considered important, but less so than human-like communication. In terms of role robots were seen by the majority as suitable for personal assistant duties carrying out household tasks. Child care or friendship roles were seen as less suitable.

Existing human-human interaction studies are a good starting point for HRI research, but can only be treated as such. Robots are not people, and not all insights and results will remain valid for HRI scenarios. So given that the nature of the interaction between humans and robots is likely to be different from that between two humans, or between humans and most current consumer technology, there are many open questions. Most importantly for the general acceptance of robots, what appearance and modalities of communication are optimal for the majority of non-technical users? Will people find a machine with a human appearance or that interacts in a human-like manner engaging or frightening? If a face is humanoid, what level of realism is optimal? What role could timing in communication [25] and the movement and timing of interactive behaviour (*kinesics* [21], [1]) play?

II. Considering Design

A. The Extended Uncanny Valley

The effect of the aesthetic design of a robot is an area that has often been neglected, and only in visual science fiction media or recently with the advent of commercial household robots has it been paid much attention. A notable exception is the 'uncanny valley' proposed by Masahiro Mori in the late 1970's [17], [5]. Mori proposed that the acceptance of a humanoid robot increases as realism increases. However there comes a point where, as the robot approaches perfect realism, the effect becomes instead very disturbing and acceptance plunges, because the robot starts to look not quite human or at worst like a moving corpse (Fig. 1). In theory the realism of both appearance and movement can give rise to this effect, with movement evoking the stronger response. It is possible that there may also be 'behavioural uncanniness' affecting perception of a robot during social interaction and governed by (among other things) the appropriateness and timing of its reponses to social cues. However little empirical data exists to support Mori's theory and opinions vary as to the strength of the effect and its longevity - our initial observations with KASPAR (section III) indicated that people soon became habituated to the robot and that feelings of uncanniness decreased rapidly with time and experience. See [15], [14] for recent work on the uncanny valley by MacDorman.

This work 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 Unit E5 (Cognition) of FP6-IST under Contract FP6-004370.

All authors are with the University of Hertfordshire, Hatfield, UK. Contact: mike@artificiallife.co.uk, K.Dautenhahn@herts.ac.uk

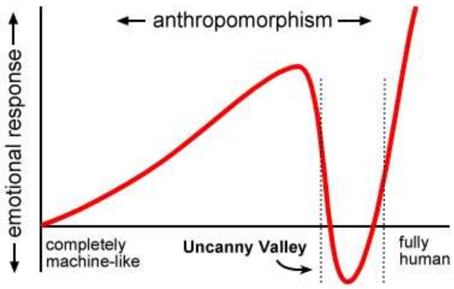

Fig. 1. Mori's uncanny valley hypothesis.

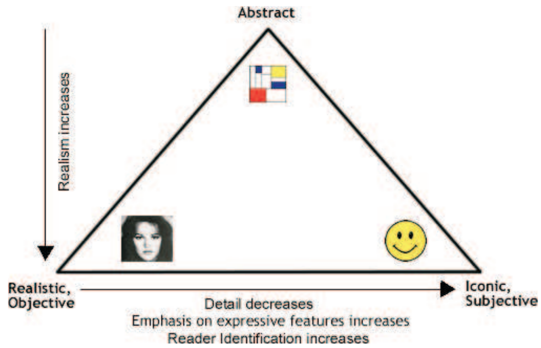

Fig. 2. The design space of faces in comics and narrative art (modified from [16]).

B. Managing Perceptions

DiSalvo *et al.* performed a study into how facial features and dimensions affect the perception of robot heads as human-like [7]. Factors that increased the perceived humanness of a robot head were a 'portrait' aspect ratio (i.e. the head is taller than it is wide), the presence of multiple facial features and specifically the presence of nose, mouth and eyelids. Heads with a landscape aspect ratio and minimal features were seen as robotic. They suggest that robot head design should balance three considerations: 'human-ness' (for intuitive social interaction), 'robot-ness' (to manage expectations of the robot's cognitive abilities) and 'productness' (so the human sees the robot as an appliance). The idea of designing a robot to be perceived as a consumer item is noteworthy for the fact that people's *a priori* knowledge of electronic appliances can be utilised in avoiding the uncanny valley; the implication is that the robot is non-threatening and under the user's control. To fulfill their design criteria they present six suggestions: a robot should have a wide head, features that dominate the face, detailed eyes, four or more features, skin or some kind of covering and an organic, curved form.

C. Faces

Faces help humans to communicate, regulate interaction, display (or betray) our emotions, elicit protective instincts, attract others, and give clues about our health. Several studies have been carried out into the attractiveness of human faces, suggesting that symmetry, youthfulness and skin condition [9] are all factors. Famously Langlois and Roggman [12] proposed that an average face - that is, a composite face made up of the arithmetic mean of several individuals' features - is fundamentally amd maximally attractive (although there are claims to the contrary, see [19]), and that attractiveness has a social effect on the way we judge and treat others [11].

Human infants seem to have a preference for faces, and it appears that even newborns possess an 'innate' ability to spot basic facial features, such as a pair of round blobs situated over a horizontal line which is characteristic of two eyes located above a mouth. It has been debated whether this is due to special face recognition capability or sensory-based preference based on preferences for general perceptual features and broad visual cues and properties of figures such as symmetry, rounded contours etc. which form the basis for learning to recognize faces [8]. The nature and development of face recognition in humans is still controversial. Interestingly, while the baby develops, its preference for certain perceptual features changes until a system develops that allows it to rapidly recognize familar human faces. Evidence suggests that exposure to faces in the first few years of life provides the necessary input to the developing face recognition system, e.g. [18]. The specific nature of the face stimuli during the first year of life appears to impact the development of the face processing system. While young infants (up to about 6 months of age) can discriminate among a variety of faces belonging to different species or races, children at around 9 months (and likewise adults) demonstrate a face-representation system that has become more restricted to familiar faces. The social environment, i.e. the 'kinds of faces' an infant is exposed to influences the child's preferences for certain faces and abilities to discriminate among them. Not only time of exposure, but also other factors, including emotional saliency, are likely to influence the tuning of the face recognition systems towards more precision [18].

D. The Design Space of Faces

In his book *Understanding Comics* [16], Scott McCloud introduces a triangular design space for cartoon faces (Fig. 2). The left apex is *realistic*, i.e. a perfect representation of reality, for example a photograph, or realistic art such as that by Ingres. Travelling to the right faces become more *iconic*, that is, the details of the face are stripped away to emphasise the expressive features; emoticons such as **:)** are a perfect example in the 21st century zeitgeist. The simplification has two effects. Firstly it allows us to amplify the meaning of the face, and to concentrate on the message rather than the medium. Secondly the more *iconic* a face appears the more people it can represent. Dautenhahn points out that iconography can aid the believability of a cartoon character [4]. We are more likely to identify with Charlie Brown than we are with Marilyn Monroe, as a realistic or known face can only represent a limited set of people whereas the iconic representation has a much broader range - to the extent of allowing us to project some of ourselves onto the character. Towards the top apex representations

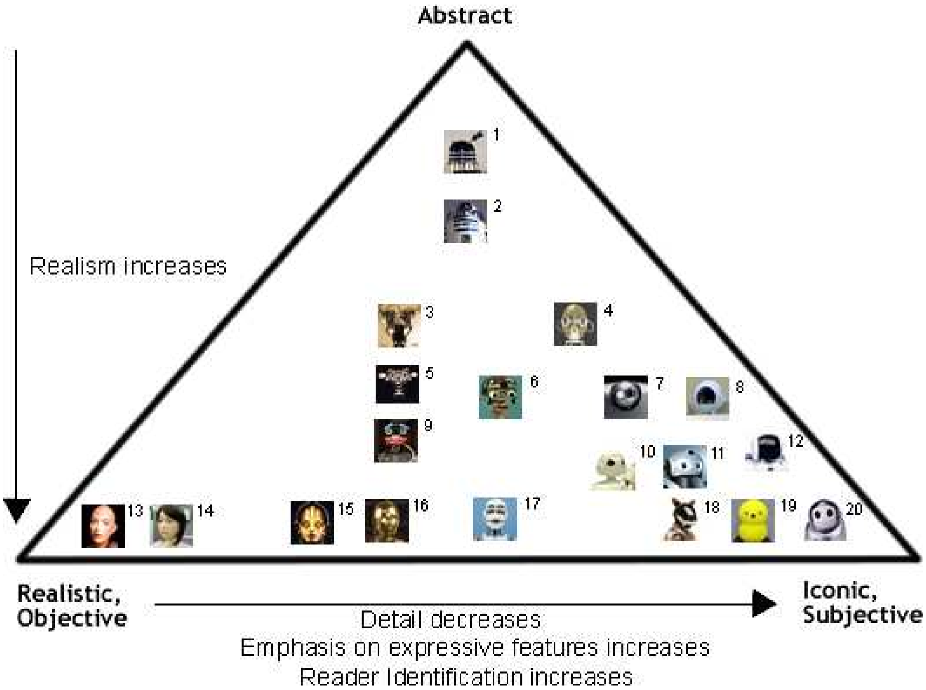

Fig. 3. Robot faces mapped into McCloud's design space. 1. Dalek ((©the British Broadcasting Corporation/Terry Nation), 2. R2D2, fictional robot from *Star Wars* (©Lucas Film Ltd.), 3. DB (©ATR Institute Kyoto), 4. MIT Humanoid Face Project (©MIT), 5. Kismet (©MIT/Cynthia Breazeal), 6. Infanoid (©Hideki Kozima), 7. Nuvo companion robot (©ZMP Inc.), 8. HOAP-2 (©Fujitsu Automation), 9. Minerva tour-guide robot (©Carnegie Mellon University), 10. Toshiba partner robot (©Toshiba), 11. QRIO (©Sony), 12. ASIMO (©Honda), 13. K-Bot, extremely realistic 24 DOF head built by David Hanson (©Human Emulation Robotics), 14. Repliee-Q1 (©Osaka University/Kokoro Inc.), 15. False Maria, fictional robot from Fritz Lang's 1927 film *Metropolis*, 16. C3PO, fictional robot from *Star Wars* (©Lucas Film Ltd.), 17. WE-4R robot (©WASEDA University), 18. AIBO robotic dog (©Sony), 19. Keepon, minimal DOF HRI robot (©Hideki Kozima), 20. Papero household robot (©NEC)

become *abstract*, where the focus of attention moves from the meaning of the representation to the representation itself. Examples in art would be (to a degree) Picasso's cubist portraits or the art of Mondrian.

E. Robot Faces in the Design Space

We can use this design space, and the accumulated knowledge of comics artists, to inform the appearance of our robots. Fig. 3 shows some robot faces and their (subjective) places on the design triangle. Most are 'real-life' robots although several fictional robots have been included, as functionality has no bearing on our classification in this context. At the three extremes are NEC's Papero (iconic), a small companion robot which is relatively simple and cheap to make and allows easy user-identification; Hanson's K-bot (realistic), complex and theoretically deep in the uncanny valley but allowing a large amount of expressive feedback, and a Dalek (abstract), potentially difficult to identify with but not as susceptible to the uncanny valley due to its non-human appearance.

Of course the design space only addresses the static appearance of the robot. The nature of most robot faces is that they encompass a set of temporal behaviours which greatly affect our perception of them. An extension of McCloud's design space to investigate behavioural aspects would be a worthwhile study, specifically how a robot's behaviour affects its perception as iconic, realistic or abstract, and the effect of social behaviour on the uncanny valley and user identification with the robot.

F. The Robot as an Extension of Self?

As one moves in the design space of the faces from realism towards iconicity, a human is more likely to identify themselves with the face due to the decrease in specific features, and the distinction between *other* and *self* becomes less and less pronounced. Could this idea be useful in robot design? If a robot is to be designed to extend the human's abilities or carry out tasks on their behalf, iconic features may possibly allow the user to more easily project their own identity onto the robot. In contrast, realistic face designs will be seen objectively as 'someone else', and *abstract* designs often as 'some*thing* else'. In this case the interaction partner's identification with the robot will be discouraged by the non-iconic nature of the design. Some robot roles (such as security guards) might benefit from reinforcing this perception. While the idea of the robot as an extension of self remains speculative at this point, future work in this area needs to shed more light on these issues.

III. KASPAR

Fig. 4 shows KASPAR (**K**inesics **A**nd **S**ynchronisation in **P**ersonal **A**ssistant **R**obotics). KASPAR is a child-sized robot which acts as a platform for HRI studies, using mainly expressions and gestures to interact with a human. The robot is a work-in-progress but when finished will comprise a static body with an 8 DOF head and two 6 DOF arms. Important features of KASPAR are minimal design, the inclusion of eyelids, and aesthetic consistency of the face (which is why eyebrows were not implemented; any mechanism to actuate them would have protruded through the skin).

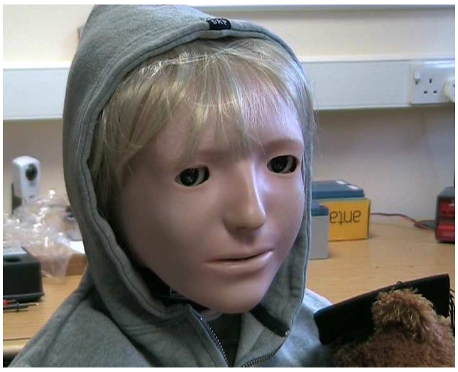

Fig. 4. KASPAR, HRI research robot.

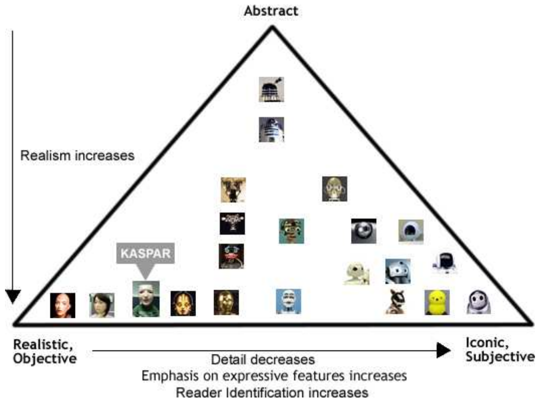

Fig. 5. KASPAR on the design space of robots.

A. Design Motivations and Rationale

Part of Hertfordshire's input in the early stages of the RobotCub project was to suggest design motivations that would help produce a useful platform for HRI studies, and which also formed the basis of the design rationale for KASPAR. These were that there should be consistency of appearance and complexity between the head, body and hands to aid natural interaction, and also between the appearance and the capabilities of the robot to govern the human's expectations. It was also suggested that minimal expressive features should be included and that they should be used to create the impression of autonomy by (for example) allowing joint attention or expressing emotional state.

The overall hardware costs of KASPAR are in the range of a desktop PC, and by keeping the complexity and DOFs to a minimum we aim to reduce building and maintainance costs while still creating a robot capable of a wide range of behaviours. The goal in this case is not perfect realism, but optimal realism for rich interaction.

B. Face Design

The face design echoes the overall rationale, in that it aims to approximate the appearance and movements of the human face without venturing into ultra-realism. Fig. 5 shows the approximate position of KASPAR on the design space of robot faces. The decision to position the face somewhat in the iconic direction was made with a two-fold purpose. We have seen that emphasis on the features used for communication allows the robot to present facial feedback clearly, by allowing the interaction partner to focus on the message more than the medium. Furthermore a reduction in detail de-personalises the face and allows us to project our own ideas on it and make it, at least partially, what we want it to be. These are both potentially desirable features for a robot in HRI scenarios. Note, however, that the emphasis on the communicative features is achieved not by using discrete, exaggerated versions (which is the case with robots such as Feelix [3] and Kismet [2]), but by reducing the distracting effect of other details of the face. KASPAR's expressions are not as unambiguously defined as those of Kismet or Feelix, but initial observations indicate that surprisingly subtle changes in expression can be effective (see experimmental results, section IV).

KASPAR's skin (a resuscitation doll mask) is only fixed at the ears and nose, and allows the face to be pulled into some fairly natural-looking expressions as the actuation of the mask in one place tends to slightly deform other areas; for instance, a smile also pushes up the cheeks and narrows the eyes. In humans this is typically considered an 'honest' smile compared to one which moves only the mouth [1].

C. Design Specifics

Requirements and Strategy. KASPAR's design was informed by initial studies of existing robot heads and by the application of ideas from McCloud's design space. The design requirements were: (1) Minimal design, yet expressive enough for HRI, (2) capacity to display *autonomy*, (3) capacity to display *undirected* and *directed attention*. (4) iconicity, (5) capacity to accept "projected" expressions with change of view angle (a requirement that was inspired by this ability in traditional Japanese noh masks [13]), and (6) human-like appearance.

Metal rods are used to transmit servo movement to the required part of the face or head. In addition to CMOS cameras in the eyes, micro-switches will be incorporated in the hands to provide simple tactile feedback and microphones added to the head.

D. Potential Uses

KASPAR can be used to study a variety of research issues relevant to HRI such as interaction dynamics, gesture creation and recognition, joint attention, communication through imitation and the use of expressions. The addition of arms will allow a range of interaction games to be played.

IV. SMILE EXPERIMENT

The first study to be undertaken with KASPAR investigated people's perception of the robot's expression. Such an experiment was considered necessary in order to provide baseline results that will inform future experiments where KASPAR's expressions will be used in regulating interaction dynamics with people. For this purpose, a simple experiment

was created to investigate what bearing the speed and continuity of a transition from one expression to another might have on the perception of a robot. Our expectations were that:
(1) Static expressions of a smile will be judged less appealing by subjects than expressions with dynamic transitions from a neutral expression.
(2) Dynamic expressions with transitions at natural speed will be judged more appealing than those with abrupt transitions.
(3) The larger the smile, the better will subjects recognize the expression of 'happiness'.

As this experiment investigates the use of movement in robot perception it can only be partially related to the idea of the design space which only concerns static images.

A. Methodology

Four degrees of smile were programmed into KASPAR and recorded on video with a plain static background. These were neutral (i.e. no smile, and the 'default' starting condition for all other expressions), and small, medium and large smiles (Fig. 6). Ten videos were created of 6 seconds duration each, showing:
1) The neutral face with no transition (static) as a control condition.
2) Small medium and large smiles with no transition (static).
3) Small, medium and large smiles with a natural transition (one that takes up to 2 seconds from neutral to smile).
4) Small, medium and large smiles with a sudden transition, created by editing the video to cut abruptly from neutral to finished smile with no intermediate stage.
It is important to note that the three sizes of the smiles remained consistent across all videos, and that only the transitions varied. A website was created which, after gathering consent and some minimal demographic data, presented all the videos twice in a random order. For each video the subject was asked to rate how happy, and how appealing, the robot's smile looked on a scale of 1-5, where 5 is maximal. Ratings of 'happiness' were expected to reflect how successful the robot's design conveyed this expression. As perceived 'happiness' could simply be interpreted as 'the amount of smile', we were also interested in how the robot would be regarded by subjects both visually and behaviourally and thus chose the term 'appealing' in an attempt to communicate the idea of this subjective judgement. All results were stored in a database for later analysis.

B. Results

Results from 51 subjects were obtained, from the UK, Norway, Sweden, Netherlands, Germany, Austria, Poland, Spain, Portugal and Italy. The subjects ranged in age between 23 and 58, 21% were female and almost all worked in a variety of academic and administration roles in universities.

1) 'Happiness' Rating: Fig. 7 shows the mean responses (average standard deviation = 0.94) to the question 'On a scale of 1-5, how HAPPY does this robot's smile look?' for each video. For the small and medium smiles, those with transitions (4-9) are perceived as marginally happier than the corresponding static smiles (1-3). For the more obvious large smile, the static version is seen as happiest followed by the natural and sudden transition versions. It is interesting that there is such a distinct classification, especially between the small and medium smiles, as at first glance the difference between them is quite subtle.

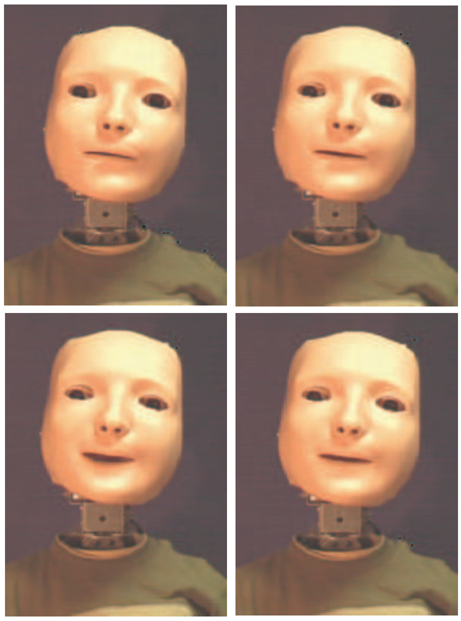

Fig. 6. The four experimental expressions, clockwise from top left: neutral, small, medium and large smiles.

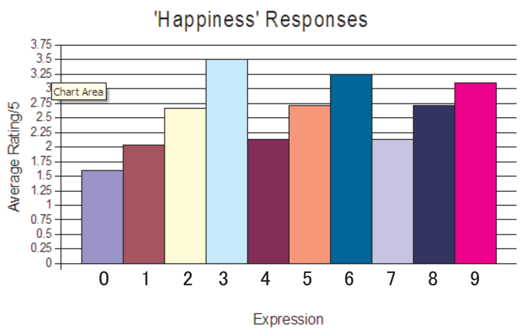

Fig. 7. Perceived 'happiness' responses. 0 = neutral, 1 = small static, 2 = medium static, 3 = large static, 4 = small NT, 5 = medium NT, 6 = large NT, 7 = small ST, 8 = medium ST, 9 = large ST. NT = natural transition, ST = sudden transition. Average standard deviation = 0.94.

2) 'Appeal' Rating: Fig. 8 shows the mean responses (average standard deviation = 1) to the question 'On a scale of 1-5, how APPEALING does this robot's smile look?' for each video. Here the clear winners are the natural transitions (4-6). In each of the small, medium and large cases the natural transition smile is rated higher than either the corresponding static or sudden transition options. Interestingly the large smile with a natural transition (6) is the most appealing of all the large smiles (in fact the most appealing of all the expressions), and yet the large smile with a sudden transition

(9) is the least. This suggests that realism or time taken to attain an expression might be a crucial factor in how the robot is perceived by human subjects.

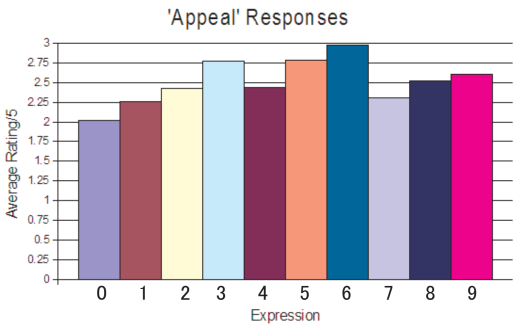

Fig. 8. Perceived 'appeal' responses. 0 = neutral, 1 = small static, 2 = medium static, 3 = large static, 4 = small NT, 5 = medium NT, 6 = large NT, 7 = small ST, 8 = medium ST, 9 = large ST. NT = natural transition, ST = sudden transition. Average standard deviation = 1.

Two of our hypotheses are supported by the results - natural transitions are seen as more appealing than sudden ones (hypothesis 2), and the larger the smile the greater the judgement of 'happiness' (hypothesis 3). However hypothesis 1 is only partially supported; smiles with a natural transition are seen as more appealing than static ones, but those with a sudden transition are not. We would suggest that in the latter case the inconsistency between appearance (fairly natural) and behaviour (unnatural) causes a negative response.

V. CONCLUSIONS

In this paper we focussed on design issues of robot faces integrating findings from psychological studies, work on narrative art design, and recent HRI studies. Consideration of these design issues strongly influenced our creation of a minimally expressive humanoid face, part of the robot KASPAR. Dimensions of face design were discussed with aims to help researchers and designers understand and exploit some ideas synthesizing those of artists, roboticists, and psychologists that pertain to human perception of robot faces in HRI. Expressions with a natural transition time are experimentally shown to be seen as more appealing than static ones or those with a sudden transition. Whether the preferred style of expression is one which has natural timing in any context or is merely one consistent with the overall aesthetic of the robot is an open question. Although these results are specific to KASPAR it is clear that robot design affects peoples' perceptions in significant ways and these results suggest that aesthetic/behavioural consistency and the temporal element in HRI are worthy of further investigation.

REFERENCES

[1] R. L. Birdwhistell. *Kinesics and Context*. University of Pennsylvania Press, Philadelphia, 1970.
[2] C. L. Breazeal. *Designing Sociable Robots*. MIT Press, 2002.
[3] L. Cañamero and J. Fredslund. I show you how I like you - Can you read it in my face? *IEEE Trans. Systems, Man & Cybernetics, Part A*, 31(5):454–459, 2001.
[4] K. Dautenhahn. Design spaces and niche spaces of believable social robots. In *Proc. IEEE Intl. Workshop Robot and Human Interactive Communication*, pages 192–197, 2002.
[5] K. Dautenhahn. Socially intelligent agents in human primate culture. In S. Payr and R. Trappl, editors, *Agent Culture: Human-Agent Interaction in a Multicultural World*, pages 45–71. Lawrence Erlbaum Associates, 2004.
[6] K. Dautenhahn, S. Woods, C. Kaouri, M. Walters, K. L. Koay, and I. Werry. What is a robot companion - friend, assistant or butler? In *Proc. IEEE IROS*, pages 1488–1493, 2005.
[7] C. DiSalvo, F. Gemperle, J. Forlizzi, and S. Kiesler. All robots are not created equal: The design and perception of humanoid robot heads. In *Proc. Designing Interactive Systems*, pages 321–326, 2002.
[8] M. J. Johnson and J. Morton. *Biology and Cognitive Development: The Case of Face Recognition.* Blackwell, 1991.
[9] B. C. Jones, A. C. Little, D. M. Burt, and D. I. Perrett. When facial attractiveness is only skin deep. *Perception*, 33(5):569 – 576, 2004.
[10] T. Kanda and H. Ishiguro. Communication robots for elementary schools. In *Proc. AISB'05 Symposium Robot Companions: Hard Problems and Open Challenges in Robot-Human Interaction*, pages 54–63, April 2005.
[11] J. Langlois, L. Kalakanis, A. Rubenstein, A. Larson, M. Hallam, and M. Smoot. Maxims or myths of beauty? A meta-analytic and theoretical review. *Psychological Bulletin*, 126:390–423, 2000.
[12] J. Langlois and L. Roggman. Attractive faces are only average. *Psychological Science*, 1:115–121, 1990.
[13] M. Lyons, R. Campbell, A. Plante, M. Coleman, M. Kamachi, and S. Akamatsu. The noh mask effect: Vertical viewpoint dependence of facial expression perception. *Proc. Royal Soc. London*, 267:2239–2245, 2000.
[14] K. F. MacDorman. Androids as an experimental apparatus: Why is there an uncanny valley and can we exploit it? In *CogSci-2005 Workshop: Toward Social Mechanisms of Android Science*, pages 106–118, 2005.
[15] K. F. MacDorman, T. Minato, M. Shimada, S. Itakura, S. Cowley, and H. Ishiguro. Assessing human likeness by eye contact in an android testbed. In *Proc. XXVII Ann. Meeting of the Cognitive Science Society*, 2005.
[16] S. McCloud. *Understanding Comics: The Invisible Art*. Harper Collins Publishers, Inc., 1993.
[17] M. Mori. Bukimi no tani [the uncanny valley]. *Energy*, 7:33–35, 1970.
[18] O. Pascalis, L. S. Scott, D. J. Kelly, R. W. Shannon, E. Nicholson, M. Coleman, and C. A. Nelson. Plasticity of face processing in infancy. *PNAS*, 102(14):5297–5300, 2005.
[19] D. Perrett, K. May, and S. Yoshikawa. Attractive characteristics of female faces: preference for non-average shape. *Nature*, 368:239–242, 1994.
[20] J. Pransky. AIBO - the No. 1 selling service robot. *Industrial Robot*, 28(1):24–26, 2001.
[21] B. Robins, K. Dautenhahn, C. L. Nehaniv, N. A. Mirza, D. François, and L. Olsson. Sustaining interaction dynamics and engagement in dyadic child-robot interaction kinesics: Lessons learnt from an exploratory study. In *Proc. 14th IEEE Ro-Man*, 2005.
[22] B. Robins, K. Dautenhahn, R. te Boekhorst, and A. Billard. Effects of repeated exposure to a humanoid robot on children with autism. In *Proc. Universal Access and Assistive Technology (CWUAAT)*, pages 225–236, 2004.
[23] G. Sandini, G. Metta, and D. Vernon. Robotcub: An open framework for research in embodied cognition. In *Proc. IEEE-RAS/RSJ Intl. Conf. Humanoid Robots*, 2004.
[24] S. Thrun, M. Bennewitz, W. Burgard, A. Cremers, F. Dellaert, D. Fox, D. Haehnel, C. Rosenberg, N. Roy, J. Schulte, and D. Schulz. Minerva: A second generation mobile tour-guide robot. In *Proc. IEEE Intl. Conf. Robotics and Automation (ICRA'99)*, 1999.
[25] M. Yamamoto and T. Watanabe. Time lag effects of utterance to communicative actions on robot-human greeting interaction. In *Proc. IEEE Intl. Workshop Robot and Human Interactive Communication*, 2003.

B.3 Detecting and Adapting to Different Styles of Play

B.3.1 On-line behaviour classification and adaptation to human-robot interaction styles.

Dorothée François, Daniel Polani, and Kerstin Dautenhahn. On-line behaviour classification and adaptation to human-robot interaction styles. Technical Report, University of Hertfordshire, School of Computer Science, College Lane, Hatfield, UK. AL10 9AB, September 2006.

On-line Behaviour Classification and Adaptation to Human-Robot Interaction Styles

Dorothée François, Daniel Polani and Kerstin Dautenhahn

Adaptive Systems Research Group, School of Computer Science, University of Hertfordshire
Hatfield, AL10 9AB
United Kingdom

{d.francois,d.polani,k.dautenhahn}@herts.ac.uk

ABSTRACT

This paper presents a proof-of-concept of a robot that is adapting its behaviour on-line, during interactions with a human according to detected play styles. The study is part of the AuRoRa project which investigates how robots may be used to help children with autism to overcome some of their impairments in social interactions. The paper motivates why adaptation is a very desirable feature of autonomous robots in human-robot interaction scenarios in general, and in autism therapy in particular. Two different play styles namely 'strong' and 'gentle' are investigated experimentally. The model relies on Self-Organizing Maps and on Fast Fourier Transform to preprocess the sensor data. First experiments were carried out which discuss the performance of the model. Related work on adaptation in socially assistive and therapeutic work are surveyed. In future work, with typically developing and autistic children, the concrete choice of the robot's behaviours will be tailored towards the childrens interests and abilities.

Keywords

Interaction styles, adaptation in interaction, behaviour classification

1. INTRODUCTION

This study is part of the AuRoRa project [1], an ongoing long-term project which investigates the potential use of robots to help children with autism to overcome some of their impairments in social interactions [7].

Children with autism have impairments in communication, social and imagination skills. Autism is a spectrum disorder and children have very different abilities and skills. In our perspective, any robotic mediated therapy therefore needs to consider the individual nature of child-robot interactions. One constraint is to make sure that the interaction between children and the robot will be 'playful' for the children (we need to consider here the notion of playfulness as it applies in autism, cf. section 2.2 below). The advantage of making the child interact with a robotic platform is to reduce the complexity of the interaction and creating a predictable environment for play to begin with, so that it can be easier for the child to feel at ease during the interaction in order to experience and understand better the interactions taking place. The premise of our work is that, progressively, the complexity of the environment can be increased if the child is making sufficient progress.

One stream of research in the Aurora project is focusing on the potential role of the robot as a mediator, i.e. as a salient object that helps children to interact with other children or adults [14, 15, 16] . In the other stream of research we focus on the robot as an autonomous toy. Here, a main objective in our research is for the robot to be able to recognize on-line the type of interaction induced by the child so that the robot can adapt to the interaction in order to behave more appropriately to the child's specific abilities and needs. At first step towards this goal, the robot should be able to maintain 'appropriate' (i.e. intermediate, balanced) levels of interaction, e.g. not too strong and not too weak. Note, we consider the child's abilities as they are expressed through interaction with the robot, resulting in different play styles. The child's therapeutic needs in this context are not addressed directly, but only indirectly by encouraging therapeutically relevant interactive behaviour involving touch [7]. The present paper presents a proof-of-concept of a robot that is adapting its behaviour on-line during interactions with the children according to detected play styles. Specifically, we show an Aibo robot that can classify specific child-robot interactions on-line, using self-organizing maps. We demonstrate how the robot can adapt its behaviour on-line to the child (i.e. to the interaction). Importantly, this work goes beyond preliminary work that classified and characterized interactions off-line, i.e. after the interactions had taken place [17, 18, 19] .

The remainder of the paper is structured as follows. Section 2 explains more precisely the motivation of this research. Section 3 characterizes the classification process. The implementation of the algorithm is described in section 4. Section 5 describes preliminary trials. Related work is discussed in more detail in section 6. Conclusions and future work close the paper.

2. MOTIVATION

2.1 Autism

Autism refers to autistic spectrum disorders which can appear at many different degrees and refer to different skills and abilities. The main impairments highlighted by the National Autistic Society are:

Impaired social interaction: Difficulties to make a sense of a relationship with others, difficulties to guess or even understand what the other's intentions, feelings and mental states are.

Impaired social communication: Difficulties with verbal and non-verbal communication (for example, difficulties to understand facial gestures).

Impaired imagination: Difficulties to have imaginative play, for example.

As a consequence of the above impairments, children often choose a world of repetitive patterns (e.g. they often play in a repetitive way).

2.2 Play

There is no precise definition about play, mostly because many fields are involved. This multidisciplinarity also results in the coexistence of various classifications of play. Among them, a classification given by Boucher [4, 5] is particularly relevant for our study in the sense that it merges the notion of exploration with the idea of social interaction.

Play is a vehicle for learning [6]. Through certain kinds of play, children can construct some understanding, in the sense of active construction of meaning. Play can thus develop skills in many fields: logical memory and abstract thought, communication skills and social skills. Moreover, it is a medium for self expression.

Children with autism have a relative potential for play but they often encounter obstacles, the causes of which are still not clear. These impairments (among them, impairments in socio emotional inter-subjectivity, in joint attention and in Theory of Mind) impair interactions in general and, more specifically, imply a lack of spontaneous and social reciprocity during play. These three impairments, in addition to the potential deficits in higher order representation may explain the difficulties encountered in pretend play. The disability in perceiving the coherence of categories and concepts can also be a reason why autistic children perceive objects in their parts and not as the whole which is part of a weak central coherence theory.

As a result, to facilitate autistic children's play with a robot, it is necessary to focus on the interaction, because interaction is decisive in the process of learning through play. If the robot is able to identify on-line the way a child interacts with the robot, then it can adapt to it more accurately. The adaptation should lead to a level of interaction encouraging the child to continue playing, and it should lead to robot's behaviours that are more appropriate to the current child's needs. For example, the robot should be able to detect forceful interaction and regulate the interaction so that the child is still engaged in the interaction but without signs of force. From this point of view, the process of adaptation would become bidirectional: firstly the robot adapts to the child and secondly, the robot may influence the child's behaviour in return.

3. CLASSIFICATION OF INTERACTION

People are used to describing an interaction verbally, by observing and listening to what constitutes the interaction. In a natural context we evaluate interaction subjectively, which means we usually don't use any objective measure to decide if e.g. an interaction is gentle or strong, repetitive or non repetitive etc. Instead, we use our own human senses and we may use as well our previous experience from similar interactions to classify and evaluate any interaction we are involved in. The challenge in this study is to classify the interaction objectively (i.e. automatically) from the robot's point of view. The interface between the child and the robot (Aibo robot) are the different sensors of the robot. This implies that we can use these quantitative measurements to evaluate, analyze and classify any interaction. Our initial idea was to run some experiments by playing with an Aibo robot according to a predefined interaction type, thus collecting all the sensor data necessary for the later analysis in order to see if and how it could be possible to match subjective human description of interaction with quantitative data. To simplify the problem, we decided to classify the interaction into two classes only: Gentle and Strong. An interaction is classified as 'gentle' if the participant is touching the robot gently, without signs of force. Note, this may also include an interaction with a child not or almost not touching the robot. On the contrary, if the participant touches the robot with signs of force, then the interaction is classified as 'strong'.

For such a classification, what is important is how a participant touches the Aibo and not which part of the robot the participant is touching. Consequently, the sensors which will contribute to the input data for the classification will be regarded as one global variable. This is possible by normalizing the input data values (repartitions into 10 bins) and computing the sum of these normalized data.

Moreover, in this first approach we wanted the classification process to be as independent from the robot's behaviour as possible. That is why we only focussed on sensors which are (almost) not influenced by the Aibo's motion and sounds it emits but are at the same time determined by an interaction with a child. Therefore, we considered as input sensors for the analysis only sensors corresponding to the touch of the head (1 sensor), the touch of the chin (1 sensor) and the touch of the back (3 sensors).

3.1 Analysis of temporal data

We conducted some preliminary experiments to get sensor data to analyze during the explorative phase of the study. The experimental setup was a participant interacting with an Aibo for around five minutes, who was asked to play during the whole session in the same way (either gently or strongly). In total, we did 6 runs, 3 with 'gentle interaction' and 3 with 'strong interaction'. The experiment involved two different adult participants, one person did one run with 'gentle' interaction and one with 'strong' interaction and the other participant did the other 4 runs. The idea of having two different participants was initially to decrease the risk of having a classification depending on the person interacting with the robot, however, one of the participants ended up doing most of the experiments. Note, this particular experiment is preliminary in nature, future work will involve a larger number of participants and experimental runs.

We analyzed the changes over time of the sum of the five external sensor data distributed into bins. Differences appeared clearly: when graphically displayed, temporal data from 'Gentle interaction' trials were made of many 'blobs' (see Fig. 1), while temporal data from 'Strong interaction' trials were mainly made of 'peaks' (see Fig. 2). Only one run with strong interaction was showing more confusing results but it was also because the participant was not interacting purely strongly during this run. We therefore excluded the results of this run for the further analysis. Given the visually different patterns, methods for automatic classification were investigated.

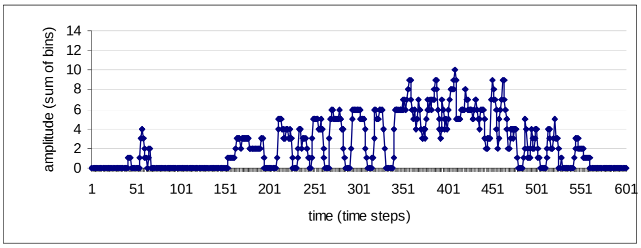

Figure 1: Gentle interaction: typical 'blobs' in temporal data.

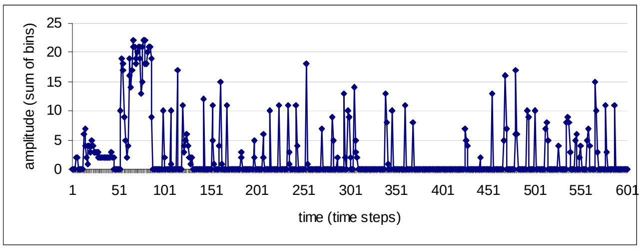

Figure 2: Strong interaction: typical 'peaks' in temporal data

3.2 Fast Fourier Transform

Since the temporal data from 'Gentle interaction' trials were made of a lot of blobs, while temporal data from 'Strong interaction' trials were mainly made of peaks it became interesting to focus on the frequency spectrum which would exhibit clear differences: for gentle interaction, there would be higher magnitudes for lower frequency and it would be the contrary for strong interaction. Moreover we wanted the method to be able to not distinguish similar patterns exhibited at different time steps: the method should be shift invariant.

Both these reasons made us select the Fourier Transform as a further step in our analysis [9]. Fourier transform is an invertible function which has, among other properties, the property of being shift invariant, and which decomposes a function into a continuous spectrum of its frequency components. Several variants coexists; among them the Fourier Transform for discrete signals and the Fast Fourier Transform which is also for discrete signals but has a complexity of $O(n \cdot \ln n)$ instead of $O(n^2)$ for the discrete Fourier Transform.

3.3 Self-Organizing Map

In a next step of the study, we wanted to automate the classification of the interaction properly, so that differences observed by eye on the magnitude of the FFT could be reflected in the quantitative analysis. Since we had no a priori information on the topology of the data, we decided to use a method which only requires poor or no a priori knowledge of the present problem and also allowed the model to learn from the data and generalize. Therefore we decided to use Artificial Neural Networks and more specifically the Self Organizing Map (SOM) which provides a topology preserving mapping from high dimensional space to map units.

SOM relies on unsupervised, competitive learning. A specific weight, from the same dimension as the input data, is attached to each neuron (node) of the network. Each node is connected to the adjacent ones according to a neighborhood rule which influences the topology of the map. The SOM is made of two phases: the training phase during which weights of the nodes are updated and the mapping phase, during which the classification or categorization of data can be made [12].

Training phase. First of all the network is initialized (either by random initialization, by initial samples, or through linear initialization). This process defines initial weight vectors, one for each node of the network. Then, input data are presented one by one to the network (random selection). For each input data, the distance is measured according to a predefined metric between the input vector and each node of the network. The node minimizing the distance is called the Best Matching Unit (BMU). Afterwards, the weights are updated according to the following equation [3] :

$w'_j = w_j + \epsilon(t) \cdot h_{rs} \cdot (v - w_j)$ $j = 1, ..., \|K\|$ where

- w_j is the weight for the node j
- w'_j is the updated weight for the node j
- $\|K\|$ is the size of neighbourhood $K(w_i(v))$ for the winner node $w_{i(v)}$.
- $h_{rs} = \exp(\frac{-d(r, w_{i(v)})^2)}{\sigma(t)^2})$ $\forall r \in K(w_{i(v)})$
- ϵ and σ are monotonic decreasing functions of time.

By simplifying, we can say that time being static, the closer a node is from the BMU, the more it will learn; and globally, the network will learn less and less when time is growing. The presentation of the entire set of input data constitutes what we call an 'epoch'. A training phase can result from the succession of many epochs.

Mapping phase. Once the network has been trained, it can be used for classifying (categorizing) data from the same space as the input data used for the training phase. A data from the latter space will be presented to the nodes successively. The node activated is the node corresponding to the BMU with regards to the same metric as used for the training phase.

3.4 Whole process of classification

The whole process of classification of interaction styles can be synthesized as follows: globally, temporal sensor data will be preprocessed to be used in the process of classification. The preprocessing results in a computation of the magnitude of the FFT algorithm which itself uses preprocessing of temporal data to consider the input sensors data as a whole global variable (see Fig. 3). The training phase of the SOM is made off-line (see Fig. 4) while the mapping phase and its coresponding preprocessing are made on-line (see Fig. 5).

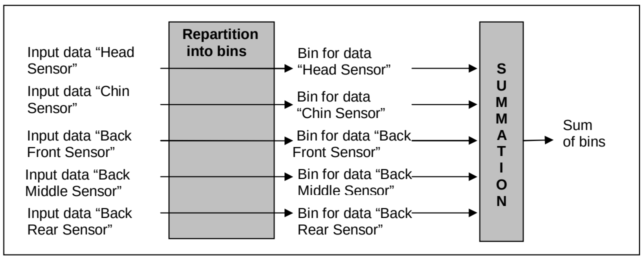

Figure 3: Data Preprocessing for FFT

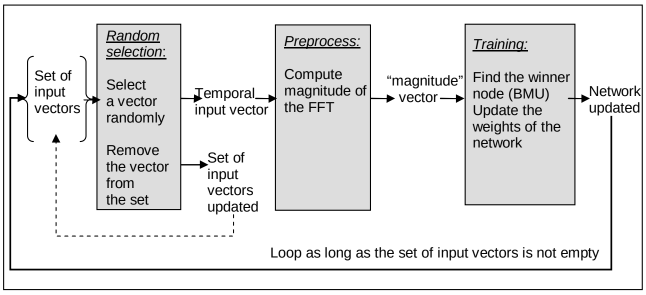

Figure 4: One Epoch of training

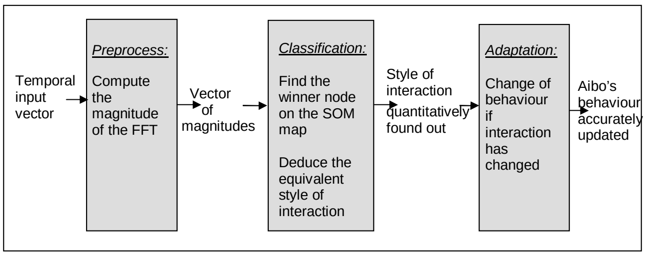

Figure 5: Unit process of classification and adaptation to the interaction style

4. IMPLEMENTATION

4.1 Communication process

The robot used in this study is the Sony Aibo ERS-7. Its control programming is achieved using URBI (Universal Real-Time Behaviour Interface) [2]. Sensor/motor data are transmitted through a wireless LAN to a laptop. The Aibo sends current states of its sensors every 32ms. The laptop analyzes periodically the sensor data, classifying on-line the interaction correspondingly and sending the information back to the Aibo which then changes its behaviour accordingly. The process of classification of the interaction is written in Java.

4.2 Parameters for the process of classification and adaptation

We needed vectors of sufficient dimension to get a good result for the SOM. Experimentally, we got good results by using an input vector of dimension 512 (it had to be a power of 2 due to the FFT algorithm we were using). Input vectors for the SOM were therefore of dimension 512 and each component of the vector was respectively the magnitude of the component of the vector resulting from the FFT. The network had a rectangular topology and was made of 10*10 nodes. We used random initialization and 5 epochs for the training which was made off-line.

However, once the training phase had been finished, all the behaviour classification was made on-line, the FFT algorithm being computed on-line as well as the activation of nodes for the SOM. But since this process was time consuming, and since the magnitude of the Fourier transform did not change significantly over a few time steps, we decided to set a frequency which would be more suitable. Experimentally it was found that updating the magnitude on the FFT once in 120 updates of the sensor data was efficient. After every update of the interaction state through the classification, the Aibo got informed of the result in order to adapt its own behaviour on-line. Note, future work will consider to run the classification on-board the robot. For monitoring and practicality purposes the use of the laptop seemed appropriate.

5. VALIDATION OF THE MODEL

5.1 Validation of the topology of the SOM map

We did two different trainings, each of them with a random intialization. We then characterized the nodes of each of the network according to the following rules: a node activated only by data from Gentle interaction is called 'gentle node'; a node activated only by data from Strong interaction is called 'strong node'; a node activated by both data is called 'hybrid node'; a node never activated is called a 'null node'.

We analyzed the topological repartition of gentle nodes on the one hand and of strong nodes on the other hand, and looked at the ratio of hybrid nodes and the ratio of null nodes. For having a performant and coherent classification of the interaction, a necessary condition is that the SOM map clearly distinguishes topologically two regions, one corresponding to the 'gentle' nodes and the second regrouping the 'strong' nodes. Moreover, the proportion of hybrid and null nodes should be very low compared to the proportion of gentle and strong nodes so that there are not two many cases in which the Aibo will not be able to 'decide' between strong and gentle interaction. Besides, hybrid nodes should be mostly on the border or next to the border between gentle and strong regions (by opposition to any of the inner part of the regions): this would correspond to a smooth transition between the two regions.

The SOM maps give both good results (see Fig. 6 which provides a graph of the first map). For each of them, the number of hybrid nodes is respectively 9 and 7 out of 100, while the number of null nodes is respectively 1 and 0. For the first map, all the hybrid nodes are on the border. For the second map, 3 hybrid nodes are not directly on the border but 2 of them are first neighbours of border nodes and the third on is second neighbour. This corresponds to a smoother transition between the two regions.

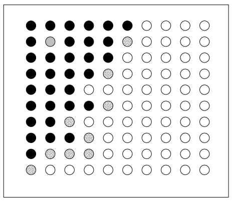

Figure 6: Map of a SOM. **legend : white for 'gentle' node, black for 'strong' node, stripes for 'hybrid' node, blobs for 'null' node**

5.2 Validation of the on-line classification and the on-line adaptation of the robot

The model is accurate if it satisfies the following constraints:

- A gentle interaction does not activate a strong node but activates one of the three other kinds of nodes; a strong interaction does not activate a gentle node but activates one of the three other kinds of nodes.
- The classification can be made on-line.
- The Aibo detects a change of the class of interaction and classifies the new interaction accurately (with an eventual short delay).
- The Aibo can adapt its behaviour on-line with respect to the kind of interaction recognized.

To test these different constraints we did various experiments with a human playing with the Aibo robot. The set of Aibo's possible behaviors remained the same in all the experiments: it was standing and waiting for at least one of its five external sensors to be activated. Whenever one of the latter sensors was activated, it started a) wagging the tail if it had detected a gentle interaction, or b) barked if it had detected a strong interaction. The rationale behind this choice was as follows: as described above, in child-robot play we want the robot to be able to maintain an intermediate level of interaction, not too strong, not too gentle. In this work, barking was used as a representative behaviour that might induce a human to 'back off', thus calming the interaction. Wagging the tail was used as an indicator to encourage interaction. Note, in future work with typically developing and autistic children the concrete choice of these behaviours will be tailored towards the children's interests and abilities.

If the Aibo had detected a middle interaction (corresponding to an activation of either null or hybrid node on the SOM map), its current reaction to tactile stimuli remained the same. Its initial state corresponded to a gentle interaction.

We ensured that the succession of interaction levels detected by the robot and the corresponding node activated on the SOM map were stored in a file. According to the experiment, the participant had to play either gently or strongly, or alternating gentle and strong interactions. The participant had to maintain the same level of interaction until the Aibo had classified and adapted to this level.

Each time the participant changed her way of interacting with the robot, the time at which it happened was stored as well as the time at which the Aibo adapted its behavior accordingly.

Note, future work will cope with more frequent changes in play style, since child users will not be instructed how to play.

Experiment 1.
In this experiment, we wanted to ensure the Aibo was able to recognize each type of interaction and keep recognizing it for the whole duration of the interaction.

This experiment consisted of two runs of three minutes each. The participant interacted with the Aibo on a gentle level of interaction only during the first run and on a strong level of interaction only during the second run. For each run, 42 updates of the classification of the interaction happened with no errors in the classification. Actually, during the 'gentle' interaction, 39 times the winner node of the SOM was a 'gentle' one, 3 times it was a 'hybrid or null' one and it was never a 'strong' one. In the same way, during 'strong' interaction, 41 times the winner node of the SOM was a 'strong' one, once it was a 'hybrid or null' one and it was never a 'gentle' one.

Experiment 2.
In this experiment, we tested the capacity of the Aibo to adapt its behaviour in a changing interaction. The participant was asked to interact gently and strongly with no constraints on the changes of the interaction styles. The only constraint was to touch quite regularly the five tactile sensors. The purpose of the experiment was to test the Aibo's capability of adaptation over time involving all five tactile sensors. We did one run that lasted around eight minutes. The Aibo adapted correctly to the interaction (see Fig. 7) but with a certain delay (which was comprised between 10s and 19s). Fig. 7 compares the Aibo's behaviour transitions (as a consequence of adaptation) to the changes in the participant's behaviour scored subjectively.

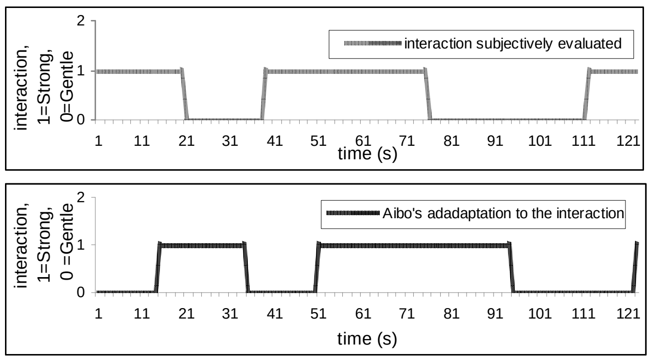

Figure 7: Example of the dynamics of the robot's adaptation to the interaction: **The first graph represents the subjective interaction level over time; the second graph shows the robot's accurate adaptation with a delay. On the y-axis, 1 stands for 'Strong' and 0 for 'Gentle'**

Experiment 3.
Since autistic children sometimes play in a repetitive way, it is very likely that some of them will continue touching the same sensor. We therefore needed to test our model in this special case and also especially because the Aibo's sensors involved in the interaction were different in nature: while the sensor on top of the head and the three sensors on the back returned values that could all vary continuously from a minimum value to a maximum value (analogical sensors), the chin sensor value could be either zero or one (numeric or boolean sensor) which means that its equivalent after repartition into bins is either 0 or bin 9.

We conducted five trials. For each of them the participant had to touch only one sensor, respectively the chin sensor, the head sensor, the back sensor on the front, the back sensor in the middle and the back sensor on the back. For each trial, the participant could change the level of interaction (from gentle to strong, from strong to gentle) whenever she wanted. Results showed that the Aibo adapted correctly to the interaction for trials focussing on the head and the three back sensors while there were some surprising results for the trial focusing on the chin sensor. We observed two kinds of possible errors in the adaptation: a) The Aibo was not able to detect a gentle interaction within 1 minute (1 minute is a long time compared to the average time of adaptation to a new interaction level), or b) The Aibo had detected a gentle interaction for a very short time (around 4 seconds), the participant was keeping interacting subjectively gently but the Aibo started barking, which means it appeared to her the interaction had become strong. This situation happened when the subjective gentle interaction was done in a way that the chin sensor was still activated (the Aibo wagged its tail). As explained above, the chin sensor can take only two values which are 0 or 1, which means, after repartition into bins, that value 1 (activation of the chin sensor) will correspond to a very high value (bin 9), even if the activation is done quite gently, (but with a sufficient pressure).

Moreover, our model for classifying data takes mainly two factors into account: a) the relative magnitude of the frequencies of the Fast Fourier Transform of one vector of sensor data exhibit which frequencies are predominant, which is directly linked to the rhythm of the interaction (e.g duration of touch of sensors, periodicity of touch of the robot on any of her five sensors etc.); b) the FFT respects the linear property. Consequently, if the chin sensor gets activated very often, even with quite gentle touch, then a lot of high values will constitute the input vector and the result of the classification may be affected.

This shows a limitation of our model: the model should be used with caution when integrating boolean sensors. If, for example, there is only one boolean sensor in five and there is a good repartition of activation of sensors, then the classification will work well. But if the boolean sensor is activated too often, it might lead to a wrong classification. Note, it also seems impossible to delimit precisely the border between strong and gentle interaction subjectively.

Experiment 4.
In the present experiment, we decided to avoid the risk of having errors induced by the boolean sensor; consequently, the participant had to respect the constraint of not touching the chin sensor, but she could touch all the four other sensors. The participant could change from one level of interaction to another (gentle, strong) whenever she wanted but she tried to vary the duration of time between the time the Aibo adapted to the current interaction and the time she changed the interaction afterwards.

The idea was to check experimentally that the delay of adaptation was not directly influenced by the rhythm of changes in the subjective interaction. This idea is linked to the fact that we use a finite vector of data to classify the interaction, which means, we take into account only a limited history of the interaction. And since this vector is updated like the process of a sliding window with a stationary length, the duration necessary to classify the interaction should belong to a very short interval of data.

The experiment lasted around seven minutes, alternating longer period for changes in behavior and shorter period for changes. The longest duration of an interaction was 50 seconds, the shortest was 17 seconds. Fig. 8. represents on the x-axis the duration of a level of interaction and on the y-axis the delay of the adaptation to the next interaction level (e.g. length of gentle interaction and delay to adapt to the next kind of interaction which will be strong). The graph shows that there is no linear relationship between the period of changes in behavior and the delay for adaptation.

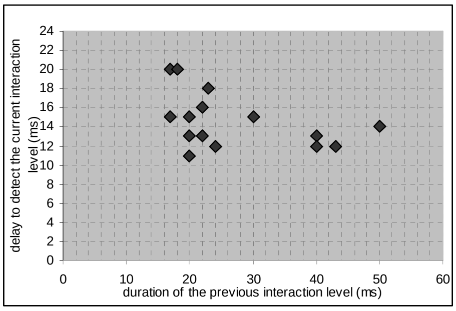

Figure 8: Experiment4 : delay in the process of classification and adaptation and corresponding duration of previous interaction style

6. RELATED WORK

6.1 Educational and therapeutic applications of human-robot interaction

Long-term (therapeutic) studies with Paro. Two studies using the seal robot Paro are particularly relevant for our study since they show that specific everyday life situations exists in which human-robot interaction can have a positive effect on well being of human beings and can even be a significiant factor of performance in therapy. The first long-term study was focusing on elderly people [22], introducing Paro into their everyday life in order to analyse the impacts on their global well being. Paro was introduced daily in two institutions for elderly people, one for 20 minutes every day over 6 weeks and the second one for 1 hour every day over more than one year. Elderly people were free to interact with Paro. Results show that the interaction with Paro improved the mood state of the participants and made the elderly people more active and more communicative with each other and with the caregivers as well.

The second study [13] designed engaging rehabilitation activities that combine physical and cognitive rehabilitation. This experiment lasted three months with a weekly occurrence. The participant was a child with severe cognitive and physical delays. The Paro robot was introduced in the Bobath protocol which is a method used for the rehabilitation of physical functional skills. Results showed that the interaction of the child with Paro seemed to have strengthened the efficacy of the Bobath protocol.

Involving quantitative data in the diagnosis of autism. The goal of this research [20, 21] is to impact the diagnosis of autism by providing the possibility to use quantitative and objective measurements of social responses. Measurements are done through both passive observation (through sensors which record and interpret data during standard clinical evaluations) and structured interactions with autonomous robots. Three criteria are mainly analyzed to distinguish typically developed children from autistic children: gaze patterns, position in the room and vocal prosody. The analysis of gaze tracking is now an integral part of the clinical evaluation. It relies on linear discriminant analysis of autistic and gaze patterns. A pilot study with this analysis has shown that autistic children don't share the same visual strategy as typically developed children and also among themselves. In this study, Scassellati exhibits a very nice application of the analysis of the interaction. He managed to qualify quantitatively criteria of typical human-human interaction through passive sensors and human-robot interaction analysis.

Long-term study on human-robot interaction in the context of dancing. This study [24, 23] aims at finding principles for realizing long-term interaction between a human and a robot. Tanaka et al. decided to run a long-term study with children and the robot QRIO, in a context relevant and frequent during childhood: dancing. This study focussed on the off-line analysis of the interaction, both qualitatively and quantitatively. On the one hand, the study analysed children's behaviour and showed that children tend to adapt their behaviour to the robot over time; e.g. they tend to know the robot is weak and tend progressively to treat QRIO softly. On the other hand, the study points out basic units as requirements for long-term interaction, respectively "sympathy" between human and robot and "variation' in the interaction style.

6.2 Classification of Human-Robot interaction

Different approaches have been used to classify human-robot interaction. More recent ones focus on the use of quantitative data for the characterisation of the interaction.

Links between subjective analysis and quantitative data. Kanda et al. [10] provide an interesting study regarding correlations between subjective evaluation (generally through questionnaires) and quantitative data collected during human-robot interaction. The experimental setup includes a participant interacting with a Robovie robot. Both are equipped with markers and infrared sources are placed in the environment. Through this setup, it is possible to collect, during the interaction, quantitative data characterizing indirectly body movements of both the robot and the subject. After the interaction phase, the individual is asked to specify the interaction subjectively according to some criteria which have been defined during a previous study [11]. The comparison between objective and subjective evaluation of the interaction indicates correlations between both analyses. In this study, Kanda et al. showed the possibility of characterizing quantitatively styles of interaction. Note, analysis of the data is off-line (i.e after the interactions have taken place) and the subjective description of the interaction focusses on the robot's behaviour only.

Salter et al. [17] adopt a different approach to show similarities between objective quantitative data and subjective description of behaviour to specify human-robot interaction. Contrary to Kanda et al.'s study, Salter et al. focus more on the participant's (a child in this study) personality during the interaction rather than on the robot's behaviour and appearance. The subjective evaluation of the children's personality takes place before the interactive phase whereby relatives of the child choose one trait of personality among a predefined list, which best corresponded to the child. The interactive phase is made of dyadic child-robot interaction with a mobile robot called Pekee (Wany Robotics); Off-line clustering analysis of the data show similarities between subjective evaluation and quantitative analysis: a) children which are considered to have the same trait of personality (among the proposed list) show also similar behaviours towards Pekee, and b) children with the same traits of personality tend to activate the same sensors on the robots (same patterns of touch).

Towards quantitative sensor analysis of the interaction. In a further study, Salter et al. [19] enumerate a list of possible states for a mobile robot called Roball and show that it is possible to define each of the states through sensor data analysis only. The four different states are: 'alone', 'interacting', 'carrying' and 'spinning'. The sensor analysis relies on off-line temporal analysis of the sensor data and a 'manual' classification through visual analysis of the sensor data which is not automated.

Automated classification and adaptation. In recent work on an adaptive playground, Derakhshan et al. [8] applied techniques known from robotics, artificial intelligence and multimedia to playgrounds. Their aim was to enable a computerized playground to adapt to children's behaviour in such a way that these children feel encouraged to play. The playground is made of specific tiles and a computer is used to store and process the data When a child is playing, input is provided through tactile sensors on the tiles. By adopting a multi-agent system approach of BDI (Belief Desire Intention) in combination with artificial neural networks techniques (with supervised training) the system learns to recognize various behaviours for either a single child or a group of children playing. Afterwards, the system can identify and adapt autonomously while children are playing. This study is very relevant to our work because it exhibits a different approach to solve the notion of on-line classification and adaptation in a context of human-computer interaction. Our study takes a different perspective though; our model aims at enabling the robotic platform to adapt its own behaviour to the interaction style, in order to a) encourage the child to continue playing, but also b) to enable the robot to influence the child's behavior to reach a specific interaction level. Note, b) is our future goal and only first steps have been taken into this direction.

7. CONCLUSION AND FUTURE WORK

This paper provided a proof of concept of on-line behaviour classification and adaptation of a robot's behaviour according to human-robot interaction styles. Experiments have shown that with our proposed model of classification a) the Aibo is able to classify a dyadic human-robot interaction it is involved in on-line, and b) it can adapt to the interaction by changing its own behaviour and thus changing the interaction with the subject.

The experiments highlighted also some limitations of the model, particularly concerning the involvement of boolean sensors in the process of collecting data. Moreover, a future step in the implementation will investigate running the algorithm on-board and will focus on an optimisation of the delay in the update of the classification of the interaction styles as well.

Concerning the process of the Aibo changing its own behaviour, more investigations needs to be done to define more accurately the different relevant behaviours for the context of child-robot interaction and more specifically for the AuRoRa project, i.e. in a therapeutic context involving autistic children. As already mentioned above, in future work with typically developped and autistic children, the concrete choice of these behaviours will be tailored towards the childrens interests and abilities.

It is hoped that this study represents a step forward in the investigation of 'child's play' with robots, involving both autistic and typically developing children.

Acknowledgments

Dorothée François is supported by a research scholarship of the University of Hertfordshire. The work described in this paper was partially conducted within the EU Integrated Project RobotCub (Robotic Open-architecture Technology for Cognition, Understanding, and Behaviours) and was partially funded by the European Commission through the E5 Unit (Cognition) of FP6-IST under Contract FP6-004370.

8. REFERENCES

[1] Aurora project. http://www.aurora-project.com.

[2] J.-C. Baillie. Urbi: Towards a universal robotic low-level programming language. In *Proc. IEEE/RSJ Int. Conf. on Intelligent Robots and Systems*, 2005.

[3] S. Baron-Cohen. *Mindblindness, an essay on Autism and Theory of Mind*. The MIT Press, new edition, 1997.

[4] J. Boucher. Editorial: interventions with children with autism - methods based on play. *Autism, The National Autistic Society*, 17, 2003.

[5] J. Boucher and P. J. Wolfberg. Editorial: Aims an design of the special issue. *Autism, The National Autistic Society*, 7, 2003.

[6] C. Chaillé and S. B. Silvern. *Understanding Through Play.* the British Library - "The World's Knowledge", 1996.

[7] K. Dautenhahn and I. Werry. Towards interactive robots in autism therapy background, motivation and challenges. *Pragmatics and Cognition*, 12, 2004.

[8] A. Derakhshan, F. Hammer, Y. Demazeau, and H. H. Lund. Mapping children playgrounds into multi-agent systems. In *Proc. 11th Int. Symposium on Artificial Life and Robotics (ISAROB)*, 2006.

[9] J. Fourier. *Théorie analytique de la chaleur*. Edition Jacques Gabay, reprinted in 1988, 1822.

[10] T. Kanda, H. Ishiguro, M. Imai, and T. Ono. Body movement analysis of human-robot interaction. In *Proc. Int. Joint Conf. on Artificial Intelligence (IJCAI)*, pages 177–182, 2003.

[11] T. Kanda, H. Ishiguro, T. Ono, M. Imai, and R. Nakatsu. Development and evaluation of an interactive humanoid robot "robovie". In *Proc. IEEE Int. Conf. on Robotics and Automation*, pages 1848–1855, 2002.

[12] T. Kohonen. *Self-Organizing Maps*. Springer, 2001 third Extended extension.

[13] P. Marti, F. Fano, V. Palma, A. Pollini, A. Rullo, and T. Shibata. My gym robot. In *Proc. AISB'05 Symposium on Robot Companion Hard Problem and Open Challenges in Human-Robot Interaction*, pages 64–73, 2005.

[14] B. Robins, K. Dautenhahn, R. te Boekhorst, and A. Billard. Robots as assistive technology - does appearance matter? In *Proc. 13th IEEE Int. Workshop on Robot and Human Interactive Communication (RO-MAN)*, 2004.

[15] B. Robins, K. Dautenhahn, R. te Boekhorst, and A. Billard. Robotic assistants in therapy and education of children with autism: Can a small humanoid robot help encourage social interaction skills? *Universal Access in the Information Society (UAIS)*, 2005.

[16] B. Robins, P. Dickerson, and K. Dautenhahn. Robots as embodied beings - interactionally sensitive body movements in interactions among autistic children and a robot. In *Proc. 14th IEEE Int. Workshop on Robot and Human Interactive Communication (RO-MAN)*, 2005.

[17] T. Salter, R. T. Boekhorst, and K. Dautenhahn. Detecting and analyzing children's play styles with autonomous mobile robots: a case study comparing observational data with sensor readings. In *Proc. of the 8th Conf. on Intelligent Autonomous Systems*, pages 61–70. IOS Press, 2004.

[18] T. Salter, K. Dautenhahn, and R. te Boekhorst. Learning about natural human-robot interaction. In *Robotics and Autonomous Systems*, 2006.

[19] T. Salter, F. Michaud, K. Dautenhahn, D. Létourneau, and S. Caron. Recognizing interaction from a robot's perspective. In *Proc. 14th IEEE Int. Workshop on Robot and Human (RO-MAN)*, pages 178–183, 2005.

[20] B. Scassellati. How social robots will help us to diagnose, treat, and understand autism. In *Proc. 12th Int. Symposium of Robotics Research (ISSR)*, 2005.

[21] B. Scassellati. Quantitative metrics of social response for autism diagnosis. In *Proc. 14th IEEE Int. Workshop on Robot and Human Interactive Communication (RO-MAN)*, 2005.

[22] T. Shibata, K. Wada, T. Saito, and K. Tanie. Human interactive robot for psychological enrichment and therapy. In *Proc. AISB'05 Symposium on Robot Companion Hard Problem and Open Challenges in Human-Robot Interaction*, pages 98–109, 2005.

[23] F. Tanaka, B. Fortenberry, K. Aisaka, and J. R. Movellan. Developing dance interaction between qrio and toddlers in a classroom environment: Plans for the first steps. In *Proc. 14th IEEE Int. Workshop on Robot and Human Interactive Communication (RO-MAN)*, 2005.

[24] F. Tanaka, J. R. Movellan, B. Fortenberry, and K. Aisaka. Daily hri evaluation at a classroom environment: Reports from dance interaction experiments. In *Proc. 1st Annual Conf. on Human-Robot Interaction (HRI)*, pages 3–9, 2006.