Learning behavior for a social interaction game with a childlike humanoid robot

Frank Broz, Hatice Kose-Bagci, Chrystopher L. Nehaniv, and Kerstin Dautenhahn
Adaptive Systems Research Group
School of Computer Science
University of Hertfordshire, UK
Email: f.broz, h.kose-bagci,c.l.nehaniv,k.dautenhahn @herts.ac.uk

I. Introduction

Social interaction plays a vital role in a child's development. At an early age, children acquire basic social skills, such as mutual gaze and turn-taking, that serve as a scaffold for more sophisticated forms of socially-mediated learning, including language. Children learn these skills through interaction with their caretakers, motivated by intrinsic social drives that cause them to seek out prolonged social engagement as a fundamentally rewarding experience. It is desirable to have robots learn to interact in a similar manner, both to gain insight into how social skills may develop and to achieve the goal of natural human-robot interaction. This paper describes a system for the learning of behavior sequences based on rewards arising from social cues, allowing the iCub, a childlike humanoid robot with a developmentally-inspired design, to engage a human participant in a social interaction game.

II. Social Cues

A. Visual Attention

Gaze is a powerful social cue. It is also one that becomes socially significant at an early developmental stage; even young infants are responsive to other's gaze direction [1]. The simplest gaze cue, and one that is the basis and developmental precursor to more complex gaze behaviors such as joint attention, is the recognition having another's visual attention [2]. This ability is crucial in a social context, as it provides valuable feedback about whether one is interacting with (or has the potential to interact with) someone or whether you are disengaged and merely sharing the same space. While there has been work done in the field of robotics using human gaze patterns to reproduce natural-appearing gaze in robots [3], [4], there has been no work explicitly modeling the role of gaze as a form of social feedback that may guide the robot's overall behavior. In this system, a gaze tracker worn by the human participant is used to collect gaze direction data in real time as sensor input for the robot (and a potential source of reward).

B. Turn-taking

Turn-taking plays a fundamental role in regulating human-human social interaction and communication whereby role-switching and the dynamics are not determined by external forces but emerge from the interaction. It has vital implications in many areas like robot-assisted therapy, especially in studies related to children with autism, where turn-taking games have been used to engage the children in social interaction [5]. Turn-taking is a skill that children begin to develop early in life. Caretakers teach infants how to engage in turn-taking through interacting with them [6]. The cues that regulate turn-taking are multimodal, and may be either general or task-based.

We hypothesize that fluid turn-taking requires attention to the recent history of both one's own and the other's actions in order to anticipate and prepare for the shift in roles. In light of this, the robot's control architecture incorporates a short term memory over the recent history of sensor data relevant to the regulation of turn-taking (to be described in Section IV). Two forms of non-verbal turn-taking are supported in this interaction, drumming and peek-a-boo. Drumming allows the human and robot to engage in turn-taking with clearly defined and easily detectable beginnings and endings. Studies on emergent turn-taking in a drumming interaction have been carried out by the authors previously using a similar childlike humanoid robot [7]. Peek-a-boo is more ambiguous (given sensing limitations) but well understood by human participants, and has also been studied before in embodied human-robot interaction [8].

III. The Interaction History Architecture

This research extends past work on the iCub using the Interaction History Architecture (IHA). IHA is a system for learning behavior sequences for interaction based on grounded sensorimotor histories. While the robot acts, it builds up a memory of past "experiences" (distributions of sensors, encoders, and internal variables based on a short-term temporal window). Each experience is associated with the action the robot was executing when it was recorded, as well as a reward value based on properties of the experience. These experiences are organized for the purpose of recall using information distance as a metric. As the robot acts, the most similar past experience to its current state is found, and new actions are probabilistically selected based on their reward value. For a full description of the architecture, see the journal article [8].

IV. Short Term Memory

In addition to a dynamic memory of sensorimotor experience and associated rewards, it is also useful to have a more detailed, fully sequential memory of very recent experience. This is especially true for skills such as turn-taking, where

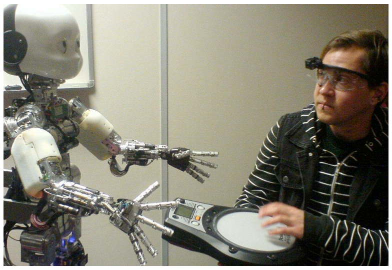

Fig. 1. A person interacts with the iCub using the gaze tracker and drum.

the recent history of relationships between one's own and another's actions must be attended and responded to. While the experience metric space preserves some ordering of experiences (so that rewards over future horizons may be computed), there is not a mechanism to recall the most recent experience, only the most similar. Additionally, experiences aggregate data over a window of time, eliminating potentially useful fine-grained information about changes in sensor values. The short term working memory preserves temporal information about the sensorimotor data over the span of several past experiences. This is especially important for guiding social interactions as it allows rewards to be designed based on these histories of interaction, rather than just the instantaneous state of the interaction that the robot is currently experiencing.

V. Developing Social Interaction

In order to demonstrate these concepts, rewards based on social drives are designed to influence the development of behaviour in an open-ended face-to-face interaction game between the iCub and a human. This work is an extension of an earlier application of IHA to the learning of the game peek-a-boo on the iCub to allow more types of interaction and social cues. The human participant interacts with the robot and may provide it with positive social feedback using their presence and gaze direction, as well as by playing a drum. The rewards representing these social drives for human presence, visual attention, and synchronized turn-taking may be based on the current state of the robot's sensorimotor experience or on the history of experience represented in its short term memory. The robot uses this feedback to acquire behavior that leads to sustained interaction with the human.

The robot comes to associate sequences of simple actions and gestures, such as waving or hitting the drum, executed under certain conditions with successful interaction based on its past experience. The pre-defined actions that the robot chooses among may be either low-level motions that don't have meaning except as part of a sequence or movements that have an (implicit) goal, such as an arm motion to hit the drum. There is no distinction in how these actions are represented internally to the architecture, and both kinds of actions may make up parts of learned behavior sequences. The extended Interaction History Architecture is intended to support the robot developing different socially communicative, scaffolded behaviours in the course of temporally extended social interactions with humans by making use of social drives and its own first-person experience of sensorimotor flow during social interaction dynamics.

VI. Future Work

The role of learning in this system is currently restricted to learning the experience space and the associations between experiences and rewards in order to find effective behavior sequences . But there are many opportunities to extend the role of learning in this system to produce social behavior in a more developmentally plausible manner. While the relationships between sensors monitored for feedback about turn-taking were predefined in this case, one could instead use statistical methods to discover which sensor channels are associated and predictive of one another. This would allow for task-specific turn-taking cues to be discovered, as well as general task-invariant cues. And while gaze is used in this system as a form of social feedback, the robot has no active gaze behavior. It would be interesting to learn action sequences for gaze behavior as well, especially gaze used to regulate turn-taking.

There is also the opportunity to engage in meta-learning about the learned behavior sequences. The topology of the experience metric space could be used to make generalizations about experiences, using the clustering of experiences to identify closely related behaviors or interactions. The ability to make aggregate representations of these clusters that capture their fundamental properties could reduce the computational cost of finding similar experiences, while possibly allowing for more powerful predictions based on current experience, opening up the potential of anticipating the actions of others or even recognizing intent.

Acknowledgment

This research was conducted within the EU Integrated Project RobotCub (Robotic Open-architecture Technology for Cognition, Understanding, and Behaviours) and was funded by the European Commission through the E5 Unit (Cognition) of FP6-IST under Contract FP6- 004370.

References

[1] S. M. Hains and D. W. Muir, "Infant sensitivity to adult eye direction." *Child development*, vol. 67, no. 5, pp. 1940–1951, October 1996.

[2] A. Frischen, A. P. Bayliss, and S. P. Tipper, "Gaze cueing of attention: visual attention, social cognition, and individual differences." *Psychological bulletin*, vol. 133, no. 4, pp. 694–724, July 2007.

[3] Y. Yoshikawa, K. Shinozawa, H. Ishiguro, N. Hagita, and T. Miyamoto, "The effects of responsive eye movement and blinking behavior in a communication robot," in *Intelligent Robots and Systems, 2006 IEEE/RSJ International Conference on*, 2006, pp. 4564–4569.

[4] B. Mutlu, J. Forlizzi, and J. Hodgins, "A storytelling robot: Modeling and evaluation of human-like gaze behavior," in *Humanoid Robots, 2006 6th IEEE-RAS International Conference on*, 2006, pp. 518–523.

[5] B. Robins, K. Dautenhahn, R. te Boekhorst, A. Billard, S. Keates, J. Clarkson, P. Langdon, and P. Robinson, "Effects of repeated exposure of a humanoid robot on children with autism," in *Designing a More Inclusive World*. Springer-Verlag, 2004, pp. 225–236.

[6] D. R. Rutter and K. Durkin, "Turn-taking in mother-infant interaction: An examination of vocalizations and gaze." *Developmental Psychology*, vol. 23, no. 1, pp. 54–61, 1987.

[7] H. Kose-Bagci, D. S. Syrdal, K. Dautenhahn, and C. L. Nehaniv, "Drum-mate: Interaction dynamics and gestures in human-humanoid drumming experiments," *Connection Science*, 2010 (In press).

[8] N. A. Mirza, C. L. Nehaniv, K. Dautenhahn, and R. te Boekhorst, "Grounded sensorimotor interaction histories in an information theoretic metric space for robot ontogeny," *Adaptive Behavior*, vol. 15, no. 2, pp. 167–187, 2007.