Socialization between toddlers and robots at an early childhood education center

Fumihide Tanaka[*†‡], **Aaron Cicourel**[§], **and Javier R. Movellan**[†]

[†]Institute for Neural Computation, University of California, San Diego, 9500 Gilman Drive, La Jolla, CA 92093-0523,[§]Department of Cognitive Science, University of California, San Diego, 9500 Gilman Drive, La Jolla, CA 92093-0515, and [‡]Information Technologies Laboratories, Sony Corporation, 5-1-12 Kitashinagawa, Shinagawaku, Tokyo 141-0001, Japan

Submitted to Proceedings of the National Academy of Sciences of the United States of America

The goal of social robotics is to develop robots that interact with people in a social manner and assist them in everyday life. Besides its technological applications, social robotics offers unique scientific opportunities to help understand the development of social interaction in humans. Here we present results of a project in which a small humanoid robot was immersed in a classroom of 18-24 month old toddlers for a period spanning more than 6 months. Three different studies are presented. In Study I we examine methods for evaluating the quality of interaction between children and robots and for testing social robot algorithms. In Study II we analyze the development of haptic behaviors during the field sessions and show that children progressively treated the robot the way they treat each other. In Study III we show that touch was a surprisingly good predictor of the perceived quality of interaction between children and robots. The study confirms that socialization and bonding between humans and robots may emerge and be sustained for significant periods of time. Haptic behaviors may play a surprisingly important role in this process.

social robotics | social development | human-robot interaction

Introduction

The development of machines capable of interacting socially with people and assisting them in everyday life has been an elusive goal of modern science [1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11]. Recent years have seen impressive advancements in the mechanical aspects of this problem, yet progress on social interaction has been slow. For example, a state of the art humanoid robot prototype, named QRIO, was built in Japan as the result of a long and costly research and development effort (Figure 1) [12, 13]. The robot displays an impressive array of mechanical skills yet its ability to interact with humans in a social manner is still almost non-existent. Research suggests that low-level information, like animacy, contingency, and visual appearance, are capable of triggering powerful social behaviors towards robots during the first few minutes of interaction [14, 15]. However it is unclear whether these social responses can be sustained for prolonged periods of time. Indeed current social robots seldom cross the “10 hour barrier”, i.e., given the opportunity, individual users spend less than a combined total of 10 hours with current robots before losing interest [1]. This is in sharp contrast, for example, with the long term interactions and bonding that commonly develops between humans and their pets. Besides its technological applications, the study of human-robot interaction offers unique scientific tools to help understand the human socialization process. Moreover, a computationally grounded theory of real-time social interaction may provide clues about the causes and potential treatments for conditions such as Autism, Asperger’s Syndrome, and William’s Syndrome.

This paper presents results from a project in which a humanoid robot was immersed in a classroom of 18 to 24 month old toddlers for a prolonged period of time. Children of this age were chosen because they have few preconceived notions of robots, and because they rely on simpler forms of social interaction that are less dependent on speech. With the exception of Study I, the robot was remotely controlled by a human operator that selected high-level behaviors, e.g., stand up and walk forward. The implementation of these behaviors into moment-by moment motor commands was handled by the robot’s own microcontrollers. The decision to run the robot in remote control mode was taken because our main goal was to understand what it takes for long-term social behaviors to emerge between humans and robots. Based on the lessons learned in this study we are currently developing algorithms for fully autonomous social interaction.

General Methods

Experimental Setup. The studies were conducted in Room-1 of the Early Childhood Education Center (ECEC) of UCSD, during the period of October 2004 to July 2005. They were approved by the UCSD IRB Board, under project #041071. Informed consent was obtained from all the parents and teachers for the studies as well as the experiments. Room-1 is divided into two indoor rooms and an outdoor playground. In all the studies, QRIO was located in the same room, and children were allowed to move freely between the different rooms (see Figure 2).

Participants. Room-1 hosts around 12 children between 10-24 months of age. In the early part of the study there were a total of 6 boys and 5 girls. In April 2005 one boy moved out and a boy and a girl moved in. The head teacher of Room-1 was assisted by two more teachers. The teachers, particularly the head teacher, were active participants in the project and provided feedback about the daily sessions.

Apparatus. QRIO is a 23 inch tall humanoid robot prototype. With the exception of Study I, it was tele-operated from a control room by the first author of this document. In addition to QRIO, two control toys were used: (1) A colorful and inviting soft toy resembling a Teddy-bear, and (2) A toy

---

Conflict of interest footnote placeholder

Insert 'This paper was submitted directly to the PNAS office.' when applicable.

[*] To whom correspondence should be addressed. E-mail: boom@mplab.ucsd.edu

[1] The 10 hour barrier was one of the concepts that emerged from the discussions at the NSF’s Animated Interfaces and Virtual Humans Workshop. Del Mar, CA, April 2004.

robot similar in appearance to QRIO. Hereafter this toy is referred to as "Robby", a name chosen by the teachers. During the 45 sessions Robby was always turned off and the children were allowed to play freely with it. All the field sessions were recorded using 2 synchronized DV cameras.

Procedure. The first and third authors of this document spent 3 months volunteering 10 hours a week at ECEC prior to the study. This allowed them to establish personal relationships with the teachers, parents, and children and helped identifying the challenges likely to be faced during the field sessions. In March 2005, QRIO was introduced to the classroom. Over time behaviors were introduced, from just sitting down on the floor to, after a week, standing up and walking around the room. Children initially were cautious about the robot and waited for about a week to start interacting actively with it. All the field sessions were conducted from 10:00 am to 11:00 am, which otherwise is commonly reserved for music and dancing activities, and lasted between 30-60 minutes. The experimental room always had a teacher as long as a child was present, and a researcher in charge of safety, often the third author of this document.

Study I: Evaluation of Social Interaction by Humans

The main goal of Study I was to develop experimental methods for scientific evaluation of social robot algorithms in the uncontrolled conditions of daily life. We evaluated two different robot-dancing algorithms: (1) A choreographed play-back dance (See Movie S6) which had been developed at great cost; (2) An algorithm in which the robot moves in response to the optic flow sensed in its cameras, resulting in behaviors that appear like spontaneous imitation dancing [17] (See Movie S7).

Methods. The study lasted for 6 sessions, 30 minutes each, and was conducted in early June 2005, by which time the children were already used to the robot. During the sessions, the robot played the same song 20 times consecutively with a 10 second mute interval before each replay. For three randomly selected sessions, the robot was controlled by the choreographed dance. For the other three sessions it was controlled by the optic-flow based dancing algorithm.

The 6 sessions were evaluated for quality of interaction by 5 undergraduate students from UCSD independently, uninformed of the purpose of the study, using continuous audience response methods originally developed for marketing research [18]: Coders operated a dial in real time while viewing the video-taped sessions. The position of this dial indicated the observer's impression of the quality of the interaction seen in the video. A computer program recorded the position of the dial and the video frame 30 times per second that the observers were seeing at that moment. Overlaid on the video, the observers could see a curve displaying their recent evaluation history (See Figure 3-A). The order of presentation of the video sessions was randomized.

Results. The evaluation signals produced by the human coders were smoothed with a low-pass filter. Figure 3-B shows the inter-observer reliability, averaged across all possible pairs of coders, as a function of the bandwidth of the low-pass filter. The inter-observer reliability shows an inverted U-curve: As the high-frequency noise components are filtered out, the inter-observer reliability increases. However, as the bandwidth of the filter decreases, it filters out more than just the noise, resulting in a deterioration of inter-observer correlation. Optimal inter-observer reliability of 0.79 (Pearson Correlation) was obtained with a bandwidth of 4.5 minutes. This suggests that the coders were implicitly averaging about 5 minutes of the past interaction when making their frame-by-frame evaluations.

Figure 3-C shows the change in the goodness of interaction score as a function of time within sessions. The dots correspond to individual sessions. The curve is the averaged score across the five judges and the 6 sessions. The graph shows a consistent decay in the evaluation score within sessions $F(1500, 7496) = 7.4768, p < 0.05$. The curve is approximately exponential with a time constant of 3.5 minutes, i.e., it takes about 4 minutes for the score to decay 36.7% of the initial value. Significant decays were also observed across sessions $F(3, 1871) = 358.07, p < 0.05$, in this case with a time constant of 1.5 sessions (i.e., it takes 1.5 sessions to reduce the score by 36.7%). The type of dancing algorithm had no significant effect, $F(1, 7496) = 2.961, p > 0.05$.

Discussion. The study showed that a simple algorithm that responds to the motion of people was as compelling as a labor-intensive choreographed dance program. It also provided parametric estimates for time constants in the expected loss of interest when performing repetitive behaviors (the quality of interaction drops by about 1/3 every 4 minutes). Most importantly the study demonstrated that experiments run in the conditions of daily life can be surprisingly efficient. While scientists agree that laboratory conditions may not be representative of daily life, it is commonly assumed that this "bias" is compensated by the reduction in variance afforded by the controlled conditions of the laboratory. However, Study I suggests that clean results may be obtained in periods of time that are shorter than those typically required for laboratory experiments. One possible explanation for this is that bringing people to the "controlled" but unnatural conditions of the laboratory may have the paradoxical effect of increasing the variance, not just the bias, of human behavior.

Study II: Emergence of Haptic Behaviors

The goal of this study was to find objective correlates of the interactions that developed between children and robots throughout the 45 field sessions. Based on extensive examination of the video tapes, we decided to focus on the analysis of haptic behaviors. Contact episodes were identified and categorized based on the part of the robot being touched: *arm/hand*, *leg/foot*, *trunk*, *head* and *face*. The coding was performed by the first author of this paper. The frame-by-frame inter-observer reliability with an independent coder uninformed of the objectives of the study was 0.85 (Pearson Correlation Coefficient).

Figure 4 shows temporally smoothed frequency counts of the different behavior categories throughout the study. With the exception of the arm/hand curve, all the other curves have an inverted U shape: Early on, when the children were cautious towards the robot, very few contact episodes were observed. Progressively, the number of times the robot was touched increased, peaked at about session 15, and then later decreased as children became more accustomed to it. Overall the most frequently touched body part was the *head*. This was due to the fact that by day 11 a simple contingency was introduced so that the robot produced a giggling sound when touched on the head (See Movie S1). The introduction of this contingency helped "break the ice" between the children and the robot, and had a dramatic effect on the improvement of the quality of the interaction.

Contrary to all the other categories, contact on the arms and hands of the robot did not follow an inverted U shape curve. Instead, its frequency increased steadily throughout the study, eventually becoming the most frequently targeted body part. This fact was very important for it was one of several indications that children were not losing interest in the robot, but rather were reorganizing the way they interacted with it. To understand the nature of this re-organization, an analysis of toddler-to-toddler contact episodes in the last two sessions was performed. First, toddler-to-toddler contact was classified as "intentional" or "incidental" (independent inter-observer reliability for this judgment was 0.95). Incidental contact occurred more or less uniformly across the body (38.4% arm/hand, 30.8% trunk, 30.8% leg/foot). However intentional peer-to-peer contact was primarily directed towards the arms and hands (52.9%) as compared to other body parts (17.6% face, 11.8% trunk, 11.8% leg/foot, 5.9% head). Thus, as Figure 4 shows, the children progressively reorganized the way they touched the robot, eventually converging to a distribution of haptic behaviors remarkably similar to the distribution of behaviors directed towards their peers as noted above.

With regard to the two control toys, the colorful Teddy-bear had elicited a large number of hugs in prior observations with children this age. Surprisingly, it was ignored throughout the study. Robby, the toy that resembled QRIO but did not move, received a great deal of attention but was treated very differently from QRIO. When children touched QRIO they did so in a very careful manner. Robby on the other hand was treated like an inanimate object or a "block", making it difficult to locate exactly where it was being touched. For this reason, haptic behaviors towards Robby and QRIO where analyzed using four new categories: *rough-housing*, *hugging*, *touching with objects* and *care-taking*.

"Rough-housing" referred to behaviors that would be considered violent if directed towards human beings. Figure 5-A shows that *rough-housing* behavior towards Robby was observed often but never observed towards QRIO. Hugging developed in distinctly different ways towards QRIO and Robby (See Figure 5-B). Robby received a surprising number of hugs from day one, yet the frequency of hugging decreased dramatically as the study progressed. The high frequency of hugging towards Robby during the early days of the study is interesting when considering that the Teddy-bear control toy, which was more "huggable" than Robby and QRIO, was never hugged. Qualitatively, it appeared as if Robby became a substitute target for behaviors originally intended for QRIO (See Movie S2) in a manner reminiscent of the displacement behaviors [19], reported by ethologists across the animal kingdom. The displacement hypothesis is based on the following facts: (1) The Teddy bear control toy, that had elicited more hugs than Robby during pilot work, was never hugged when QRIO was present. Robby, on the other hand was hugged an unusual amount of time when QRIO was present. (2) As the study progressed hugging towards QRIO increased while hugging towards Robby decreased dramatically. Indeed it took nearly a month for children to start hugging QRIO (See Movie S3) but once they did so its frequency was sustained until the end of the study. It should be noted that the category "hugging" included behaviors like 'holding' or 'lifting up' which were in general far more difficult to do with QRIO than Robby, which is lighter and does not move autonomously. In spite of this, by the end of the study the least huggable entity (i.e., QRIO) was hugged the most, followed by Robby. The most huggable toy, i.e., the Teddy bear, was never hugged.

Another behavioral category that developed very differently towards Robby than QRIO was "*touching with objects*". This generally involved social games that the children played with QRIO (e.g., giving QRIO an object or putting a hat on it). These behaviors were seldom directed towards Robby, but commonly occurred throughout the 45 sessions with QRIO (See Figure 5-C). *Care-taking* behaviors were also frequently observed towards QRIO but seldom towards Robby. The most common behaviors from this category involved putting a blanket on QRIO/Robby while saying "night-night" (See Movie S4). This often occurred at the end of the session when QRIO laid down on the floor as its batteries were running out.

Occasionally QRIO and Robby fell down due to contact from the children. The frequency with which this happened followed noticeably different patterns for QRIO and Robby (See Figure 5-E). Robby fell regularly throughout the duration of the study. During the first sessions QRIO seldom fell, likely due to the children's reluctance to touch it. As interactions between children and QRIO developed, the number of times it fell increased. However by the end of the study QRIO hardly ever fell despite a marked increase in the degree of interaction between the children and QRIO. Early in the study, some children cried when QRIO fell. A month into the study children seldom cried and instead they helped QRIO stand up by pushing its back or pulling its hand (See Movie S5). Overall, as Figure 5, shows, social behaviors towards QRIO increased across the entire duration of the study, while social behaviors towards Robby decreased over time.

Study III: Automatic Assessment of Social Connectedness

In Study I we developed methods for evaluating the subjective quality of interaction between children and robots. In Study II we analyzed the development of haptic behaviors. In Study III we focused on analyzing the relationship between subjective quality of interaction and objective haptic behaviors.

Methods. 15 sessions were randomly selected out of the 45 sessions, and all the videos were coded by four undergraduate students from UCSD independently, uninformed of the purposes of the study, using the same procedure we explained in Study I. A variety of statistical models were developed and tested in an attempt to predict the human evaluation of the quality of interaction. For every video frame in the 15 field sessions, the models were given 8 binary inputs indicating the presence or absence of 8 haptic behavior categories described in Study II: *touch head, face, trunk, arm/hand, leg/foot, hugging, touching with objects, and care-taking.* The goal of the models was to predict, frame by frame, the goodness of interaction score, averaged across the 4 humans coders. Amongst the models evaluated, the simplest and most successful one was structured as follows: First, the outputs of the 8 haptic input signals were low pass filtered and followed by a time delay. The time-delayed and low pass filtered signal was then linearly scaled to predict the quality of interaction averaged across the 4 human observers. Four parameters were optimized: (1) The bandwidth of the low-pass filters; (2) The time delay; (3) The additive; and (4) multiplicative constants of the linear transformation. The optimal bandwidth was 0.0033 Hz and the optimal time delay was 3 seconds. With these parameters the Pearson correlation coefficient between the model and the human evaluation of the quality of interaction, across a total 1,244,224 frames, was 0.78, almost as good as the average human-to-human agreement (0.79). More complex models were also tested that assigned different filters and different weights to different haptic behaviors, but the improvements achieved by such models were small. Figure 6 displays the evaluation of the 4 human coders and the predictions based on the touch model for a single session. Representative images are also displayed from different parts of the session.

Discussion. Human perception of the ongoing quality of interaction was predicted quite well from haptic behaviors, i.e., touch. The result is reminiscent of Harlow's famed experiments with infant macaques raised by artificial surrogate mothers. Based on those experiments, Harlow concluded that "contact comfort is a variable of overwhelming importance in the development of affectional response" and hypothesized that "contact comfort has long served the animal kingdom as a motivating agent for affectional responses" [20].

It is surprising that such a simple model could do so well. It appears that touch may be a very good indicator of social connectedness in early social interaction. There may be computational reasons for this to be the case. For example, machine perception research shows that it takes complex computations to process images and sounds in the uncontrolled conditions of daily life. Touch, on the other hand, may be easily detected and categorized [21, 22]. If so, it would be advantageous for haptic signals to be reserved for the most critical behavioral processes, e.g., to provide a sense of well being and social connectedness that could be used to bootstrap the development of social skills. Should this be the case, touch may also prove to be a useful signal for robots that learn to interact with people on their own.

Conclusions

After 45 days of immersion in a child-care center, we found evidence that long term bonding and socialization occurred between children and a robot remotely controlled by a human. Indeed rather than becoming bored, the children progressively treated the robot the way they treat each other. The results highlighted the particularly important role that haptic behaviors play in the socialization process: (1) The introduction of a simple touch-based contingency had a breakthrough effect in the development of social behaviors towards the robot. (2) The distribution of touch behaviors towards the robot converged, as the study progressed, to the distribution of touch behaviors towards other peers. (3) Touch, when integrated over a few minutes, was a surprisingly good predictor of the ongoing quality of interaction between the children and the robot. Based on these results, we advanced the hypothesis that in early social development touch may also be a very good indicator of human social connectedness, and that the human brain may use it as a signal to evaluate the quality of the ongoing social interaction. One prediction from such hypotheses is that it may be possible to develop social robots that capitalize on touch as a reinforcement signal to learn to interact with others. Another prediction is the existence of brain systems that keep track of the ongoing rate of touch at the time scale of minutes. Such a hypothesis could be tested with current brain imaging methods.

The study presented here builds upon prior work on social robotics [1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11] and suggests important lessons for progress: (1) It is possible and efficient to run experiments in the supposedly chaotic conditions of daily life. (2) It is possible for humans to socialize and bond with robots beyond the "10 hour barrier". (3) Touch may be a critical factor in the development of this socialization process.

We thank the director of ECEC Kathryn Owen, the head teacher of Room-1 Lydia Morrison, and the parents and children of Room-1 for their support. The study is funded by the UC Discovery Grant #10202 and by the NSF Science of Learning Center grant SBE-0542013.

1. Picard, R.W. (1997) *Affective Computing*, (The MIT Press).
2. Brooks, R.A., Breazeal, C., Marjanovic, M., Scassellati, B., and Williamson, M.M. (1999) *Lecture Notes in Artificial Intelligence* **1562** 52-87.
3. Weng, J., McClelland, J., Pentland, A., Sporns, O., Stockman, I., Sur, M., and Thelen, E. (2000) *Science* **291** 599-600.
4. Breazeal, C.L. (2002) *Designing Sociable Robots*, (The MIT Press).
5. Fong, T., Nourbakhsh, I., and Dautenhahn, K. (2003) *Robotics and Autonomous Systems* **42** 143-166.
6. Kanda, T., Hirano, T., Eaton, D., and Ishiguro, H. (2004) *Human-Computer Interaction* **19** 61-84.
7. Pentland, A. (2005) *IEEE Computer* **38** 33-40.
8. Robins, B., Dautenhahn, K., Boekhorst, R.T., and Billard, A. (2005) *Universal Access in the Information Society* **4** 105-120.
9. Kozima, H., Nakagawa, C., and Yasuda, Y. (2005) *Proceedings of the 2005 IEEE International Workshop on Robot and Human Interactive Communication*, 341-346.
10. Kahn Jr., P.H., Friedman, B., Perez-Granados, D.R., and Freier, N.G. (2006) *Interaction Studies* **7** 405-436.
11. Nagai, Y., Asada, M., and Hosoda, K. (2006) *Advanced Robotics* **20** 1165-1181.
12. Kuroki, Y., Fukushima, T., Nagasaka, K., Moridaira, T., Doi, T.T., and Yamaguchi, J. (2003) *Proceedings of the 2003 IEEE International Workshop on Robot and Human Interactive Communication*, 303-308.
13. Ishida, T., Kuroki, Y., and Yamaguchi, J. (2003) *Proceedings of the 2003 IEEE International Workshop on Robot and Human Interactive Communication*, 297-302.
14. Movellan, J.R., and Watson, J. (1987) *Infant Behavior and Development.*
15. Johnson, S., Slaughter, V., and Carey, S. (1998) *Developmental Science* **1** 233-238.
16. Tanaka, F., Movellan, J.R., Fortenberry, B., and Aisaka, K. (2006) *Proceedings of the 1st Annual Conference on Human-Robot Interaction*, 3-9.
17. Tanaka, F., and Suzuki, H. (2004) *Proceedings of the 2004 IEEE International Workshop on Robot and Human Interactive Communication*, 419-424.
18. Fenwick, I., and Rice, M.D. (1991) *Journal of Advertising Research* **31** 23-29.
19. Eibl-Eibesfeldt, I. (1958) *The Quarterly Review of Biology* **33** 181-211.
20. Harlow, H.F. (1958) *American Psychologist* **13** 573-685.
21. Tajima, R., Kagami, S., Inaba, M., and Inoue, H. (2002) *Advanced Robotics* **16** 381-397.
22. Miyashita, T., Tajika, T., Ishiguro, H., Kogure, K., and Hagita, N. (2005) *Proceedings of the 12th International Symposium of Robotics Research*

Fig. 1. QRIO, a small humanoid robot prototype was used as the research platform.

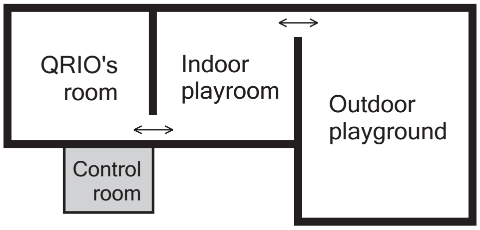

Fig. 2. Layout of Room-1 at ECEC where QRIO was immersed. There were three playing spaces, in one which QRIO was placed. Children were free to move back and forth between spaces thus providing information about their preferences.

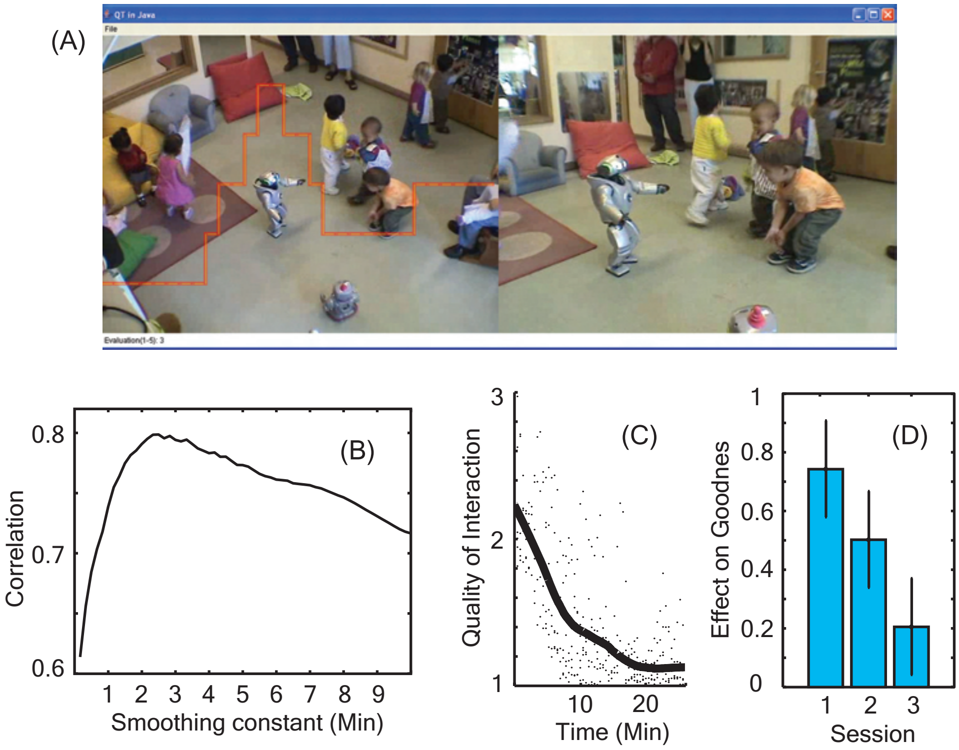

Fig. 3. (A): Each coder was asked to move a dial in continuous time to indicate their perception of the goodness of the interaction between children and QRIO observed in the video. (B): Inter-Observer reliability between four coders as a function of low-pass filter smoothing constant. (C),(D): Main effects in the goodness of interaction score as a function of time within a session (C) and across sessions (D).

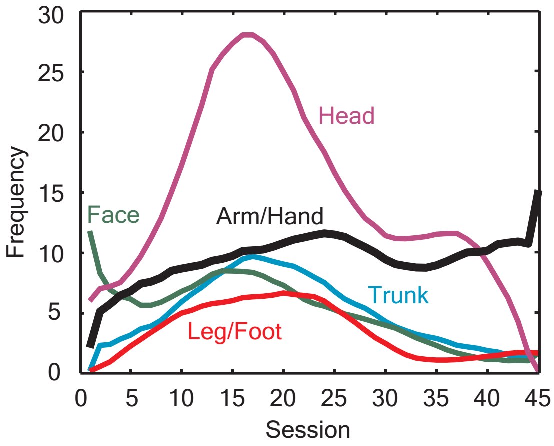

Fig. 4. Evolution of the frequency distribution of children's touch on five areas of QRIO's body over 45 sessions.

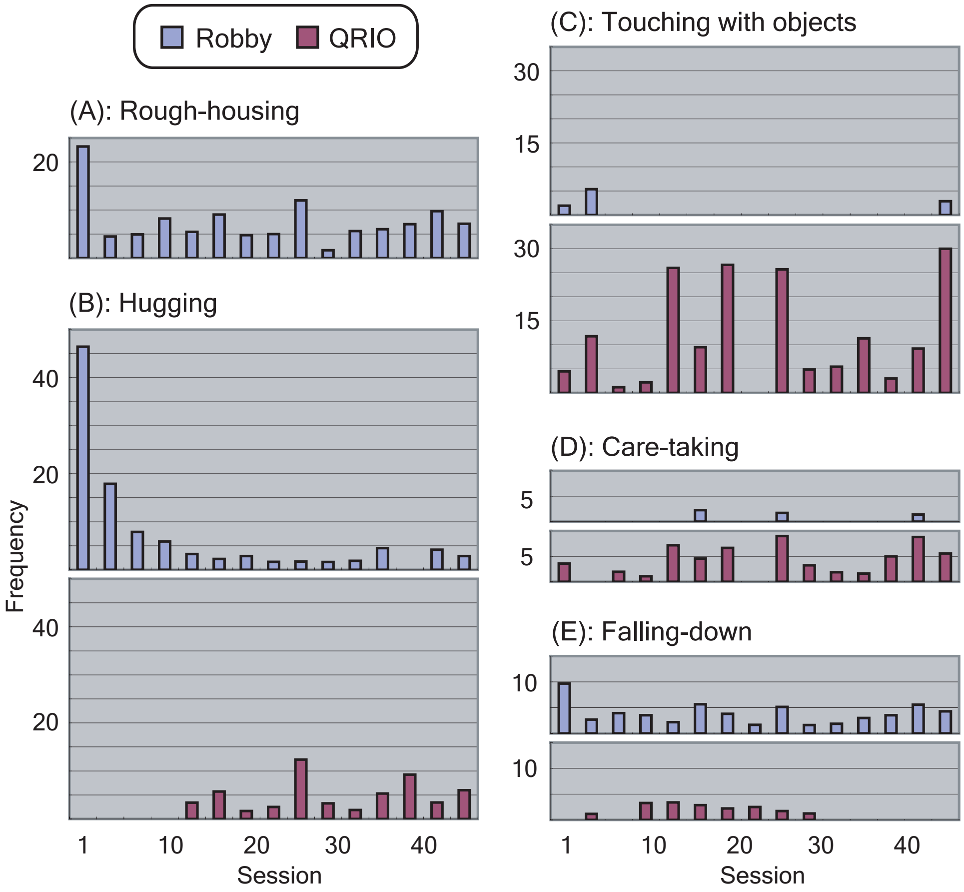

Fig. 5. Evolution of the frequency counts of different behavior categories throughout 45 daily sessions. Rough-housing was never observed towards QRIO.

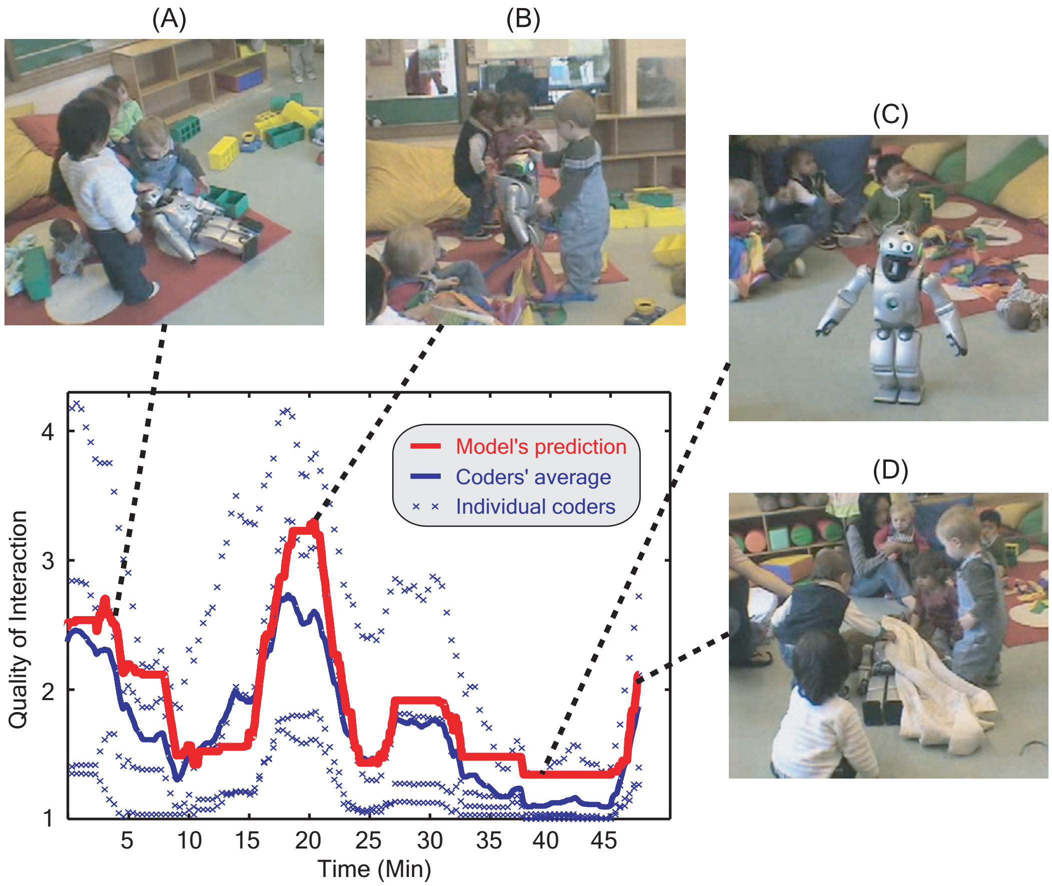

Fig. 6. Red line: Automatic assessment of the goodness of interaction between children and QRIO based on haptic sensing. Blue lines: Human (four independent coders) assessment of the goodness of interaction by using the continuous audience response method. (A): A session begins with QRIO waking up attracting children's interests. (B): During the music time in the classroom children play with QRIO. (C): Children are getting tired of the music time and loosing interests on QRIO. (D): Children put a blanket on QRIO lying down on the floor preparing for the end of a session.