The RUBI Project: A Progress Report.

Javier R. Movellan
University of California
San Diego

Fumihide Tanaka
Sony Corporation

Ian R. Fasel
University of California
San Diego

Cynthia Taylor
University of California
San Diego

Paul Ruvolo
University of California
San Diego

Micah Eckhardt
University of California
San Diego

ABSTRACT

The goal of the RUBI project is to accelerate progress in the development of social robots by addressing the problem at multiple levels, including the development of a scientific agenda, research methods, formal approaches, software, and hardware. The project is based on the idea that progress will go hand-in-hand with the emergence of a new scientific discipline that focuses on understanding the organization of adaptive behavior in real-time within the environments in which organisms operate. As such, the RUBI project emphasizes the process of *design by immersion*, i.e., embedding scientists, engineers and robots in everyday life environments so as to have these environments shape the hardware, software, and scientific questions as early as possible in the development process. The focus of the project so far has been on social robots that interact with 18 to 24 month old toddlers as part of their daily activities at the Early Childhood Education Center at the University of California, San Diego. In this document we present an overall assessment of the lessons and progress through year two of the project.

Categories and Subject Descriptors

I.2 [**Artificial Intelligence**]: Robotics

General Terms

Design,Experimentation,Human Factors

Keywords

Design by Immersion, Field Studies, Social Robots, Architectures for Social Interaction

1. PHILOSOPHY OF THE RUBI PROJECT

The development of social robots brings a wealth of scientific questions and technological challenges that are only starting to be addressed in a coordinated manner [35, 10, 48, 9, 17, 20, 33, 38, 23, 34, 32]. Progress is likely to be facilitated by a Kuhnian-style "Scientific Revolution" [24], e.g., a shift in focus and methods towards the computational analysis of real-time social interaction in everyday environments. With this idea in mind, two years ago we started a project, named RUBI, to pursue three key insights: (1) Progress can be accelerated by developing robotic systems *immersed* in the environment in which they need to operate. This allows feedback to occur on a daily basis, facilitating the rapid discovery of problems that need to be solved. (2) Social Robotics needs to emerge as a coherent program that cuts across scientific disciplines. For example, decoupling the study of human robot-interaction from computational analysis and hardware development may be detrimental in the long run. (3) Progress requires a mathematical framework for the analysis of real time social interaction.

Following Solomon Asch's views on the emergence of new scientific disciplines [2], we avoided rushing into controlled conditions and laboratory experiments. Instead our focus was on observation and identification of scientific questions and technological challenges. As such we emphasized the process of *design by immersion*. The goal was to have real life environments shape the hardware, software, and scientific goals as early as possible in the development process. We were also particularly inspired by David Marr's philosophy of science [27] which emphasizes the importance of computational analysis, i.e., understanding the computational nature of the problems that organisms solve when operating in their daily environments.

We decided to focus the project on the development of robot technologies to interact with toddlers (18-24 months of age). Children at this age were chosen because they have few preconceived notions of robots and because we believed they would help focus our work on problems we deemed particularly important, e.g., timing, non-verbal communication, and the development of affective and social bonding behaviors.

In its two years of life the RUBI project has generated a wealth of scientific articles spanning behavioral analysis of human-robot interaction [42, 43, 44, 22, 14], new machine perception primitives [6, 41, 26, 5, 4, 40], new machine learning algorithms [11, 29], and documentation of the process of *design by immersion* [30, 18]. The specifics of the discoveries emerging from the project can be found in these articles. Here we focus on a general overview of the progress and experiences accumulated during the first two years of the

project. First we describe the origins and time-line of the project. Second we summarize the progress and research activities for Years 1 and 2. Finally we provide an overview of our general experience in the project and the difficulties we found implementing the design by immersion approach.

2. ORIGINS AND TIMELINE

The conceptual seeds of the RUBI project date back to an NSF-ITR grant to develop machine perception primitives, e.g., expression recognition, for autonomous tutoring systems [12]. We realized that while such perception primitives can be developed using current technology, the problem of how to connect perception and action to produce fluid social behavior in real time is far less understood. The RUBI project started in September 2004 [30] with an aim towards studying this problem in a comprehensive way. The project is still evolving and is currently operating under the auspices of the UC Discovery Program. Below we describe the main research activities and results obtained during Years 1 and 2.

3. SUMMARY OF YEAR 1: DEVELOPMENT AND DATA GATHERING

The field studies in the RUBI project are being conducted at the UCSD Early Childhood Education Center (ECEC). The first 6 months of the project were spent volunteering at the ECEC. This time was important for bonding with the children, teachers, and parents, and for developing a sense of the problems that were likely to be encountered. It also helped shape the general philosophy of the RUBI project, as expressed in the previous sections.

The next 6 months were dedicated to conducting field sessions with two robots: RUBI, and QRIO. RUBI is a robot platform that is being designed from the ground up by immersion in the classroom. QRIO is a state of the art humanoid robot prototype developed by Sony Corporation [1, 25, 19]. During these 6 months we conducted in one room of ECEC a total of 60 field sessions (See Figure 2-Top) . All the sessions were taped with two synchronized cameras for further analysis.

4. SUMMARY OF YEAR 2

In Year 2 the focus was on analysis of the 60 videotaped field sessions and on redesigning the software for the RUBI robot prototype based on the lessons we learned as part of Year 1.

4.1 Analysis of the Field Studies

Developing efficient methods to analyze the field sessions in a manner that suited the goals of the project was not a trivial task. Over time we found two methods particularly useful: (1) The continuous audience response methods used in marketing research [36, 16], and (2) Frame-by frame labeling for the Presence/Absence of target behaviors.

Regarding the continuous audience response methods, we developed software that allowed observers to operate 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 interaction seen in the video. 30 times per second the program recorded the position of the dial and

Figure 1: *A toddler playing with QRIO (left) after a teaching session with RUBI (Right).*

the video frame that the observers were viewing at that moment. Overlaid on the video, the observers could see a curve displaying their recent evaluation history (See Figure 2-Bottom). We found that in spite of the abstract nature of the dimension being coded (quality of interaction) interobserver reliability was quite high (average Pearson Correlation between 5 independent observers was 0.79) [44, 42]. We also coded, frame by frame, the presence or absence of a variety of objective behaviors, e.g., "QRIO was touched on the head". Time series analysis revealed that haptic behaviors were surprisingly effective predictors of the perceived quality of interaction. A linear combination of the output of low-pass filtered touch sensors could predict the frame by frame quality of interaction, as assessed by humans, very well [14]. While the interpretation of this result is still unclear it helped raise our awareness about the special role that touch and haptic behaviors play in social interaction.

We also analyzed the results of an experiment conducted in Year 1. The goal of the experiment was to evaluate two different dancing algorithms for social robots [44, 42]. The study lasted 6 field sessions at ECEC, 30 minutes each. For three randomly selected sessions, QRIO was controlled by a choreographed dance program. For the other three sessions it was controlled by an optic-flow based dancing algorithm [45]. The study showed that a simple algorithm that responds to the motion of people was as compelling as a labor-intensive choreographed dance program. Most importantly, it taught us that it is possible to run experiments, not just observational studies, in the relatively uncontrolled conditions of daily life. We obtained replicable results in periods of time that were shorter than those typically required for laboratory studies. In fact we feel strongly that the experiment would have been very difficult to conduct in a laboratory setting detached from the daily conditions and activities of the children.
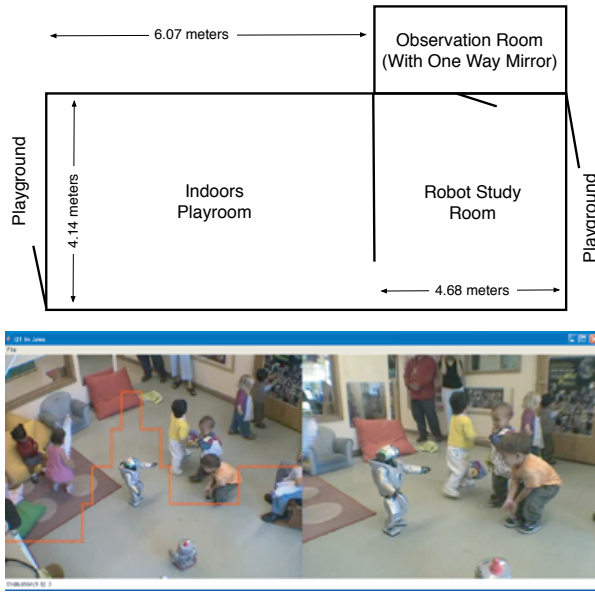
Figure 2: Top: *Layout of Room 1 at ECEC. There are three playing spaces. Children are free to move back and forth between spaces thus providing information about their preferences.* **Bottom:** *Judges used a dial to continuously evaluate the quality of the interaction between children and robots. Judges can simultaneously see two synchronized movies taken by two separate cameras. They can also see the recent history of their evaluation which is superimposed on the movie as a red graph.*

4.2 Mathematical Formalization

A critical challenge in Social Robotics will be the development of a mathematical framework for formalizing how to connect perception and action in the context of real time social interaction. As part of the RUBI project we took significant steps towards such a goal. The framework we are pursuing is based on the *Theory of Stochastic Optimal Control*, an area of mathematics and engineering that deals with the control of probabilistic processes [7]. It focuses on the problem of finding *control laws*, i.e., moment-to-moment mappings between perceptions, internal states, and actions, to achieve long-term goals. The theory of stochastic optimal control was designed to solve many of the problems that have been elusive to traditional cognitive and symbolic AI approaches, particularly in regard to social interaction: (1) the importance of timing; (2) the fact that social interaction is a continuous "dance" rather than a turn-taking process; (3) the need to act intelligently in the presence of constantly changing and uncertain information. The human brain faces many control problems when sending motor commands to the body's actuators. Riding a bicycle, using a computer mouse, shooting baskets, and playing music are all control problems. We and others believe that real time social interaction is also, in essence, a control problem. The parameters of the social interaction problem are different from the parameters of the physical interaction problem but the mathematical structure of these two problems is identical [31, 49].

We developed an example of how stochastic optimal con-

trol can be used to formalize real time social behavior. In particular we focused on the problem faced by two month old infants, given their limited perceptual capabilities, of detecting the presence of responsive human caregivers [29]. The idea behind the approach was that humans can be identified by the temporal characteristics of their response to others, a source of information commonly known as "social contingency" in infancy learning literature [3, 8, 46, 47]. From this point of view the problem faced by infants is that of detecting the "social contingency signature" hidden in the stream of activity continuously received by their sensors. Once the problem was formalized this way, an optimal controller was developed that scheduled simple vocalizations, moment to moment, so as to detect social contingency as quickly and as accurately as possible. The optimal controller exhibited some interesting properties: (1) It modeled well the temporal dynamics of vocalizations found in social contingency experiments with infants [37, 31]. (2) Turn taking behaviors emerged in the controller as an optimal strategy. Most importantly these "turns" were not fixed apriori. Its length, for example, changed dynamically based on the incoming sensory information. The algorithm was implemented in a humanoid robot and used as a primitive that allowed it to learn on its own how to detect human faces [11].

Figure 3: *Three real time control problems: person controlling a computer mouse, infant playing smile games with Mom, RoboVie-I playing with a person.*

5. SOFTWARE DEVELOPMENT

The design by immersion process was particularly influential in two aspects of software development: (1) The de-

velopment of a new software architecture for social robotics. (2) The development of machine learning and machine perception primitives robust enough to operate in field conditions.

5.1 Software Architecture: RUBIOS

We developed the first version of RUBIOS, a software architecture inspired on the ideas of stochastic optimal control. The architecture is designed to handle timing, uncertainty and learning in a network of goal-oriented message-passing nodes. The ultimate goal in RUBIOS is to have the programmer focus on the goals the robot and let probability theory and machine learning take care of the details of how to best achieve those goals. Each node in a RUBIOS robot implements a parameterized "control law", i.e., a function that couples the history of past sensor information and internal representations to current actuator outputs. In addition each node has a learning processes whose role is to change the node's parameters to maximize the long-term pay-off accumulated by that node.

Nodes can affect each other by offering rewards that vary as a function of dimensions such as time or similarity between the desired goal and alternative goals. For example, a node may offer a reward for positioning a perceived face as close as possible to a desired location on the retina. In the default implementation rewards vary as a function of time and space in an exponential/quadratic manner:

$$\rho(t,x) = \rho_0 e^{-\frac{(x-\xi)^2}{\beta}} \; e^{-\frac{t}{\alpha}} \qquad (1)$$

where ρ_0 is the peak reward, α the time scale, β the space scale, and ξ the desired position. Here "space" refers to an abstract similarity space between states. Within the RUBIOS framework reflexive processes correspond to offers with high peak value and short time constants. Moods on the other hand emerge as a result of offers with small peak values but very long time constants.

The goal of each RUBIOS node is the long term maximization of rewards. To do so each node has a simple control law, which typically involves a greedy controller, and a learning process that is in charge of optimizing the default controller. Nodes can pass information to each other via multiple channels, which are typically optimized for the type of message being passed. For example, images are passed via memory mapping, while low-bandwidth messages are typically passed via sockets. The current version of RUBIOS consists of a set of classes specialized on different aspects of robot operations: Inter-node communications, Interface with human operator, Node monitoring, Servo Control, Game control, and Vision.

While our experience with the current version of RUBIOS is still very limited, it holds promise in terms of the ease with which different programmers can seamlessly add nodes that integrate with the overall robot behavior. For example, energy saving was implemented by adding a constant request for all the servos to move to their resting point, yet having that request have a very small peak value and a very long time constant. Adaptive reflexes were also easily implemented by adding nodes that produce requests with very large peak values but very short time scales.

5.2 Perceptual Primitives

For the past 15 years our laboratory has focused on the development of perceptual primitives for social interaction (e.g., face detection, expression recognition, audio visual speech recognition). Over the years these systems have been refined and are now operating in or near real time. RUBI's software include the latest versions of face detection, face tracking and video-based emotion recognition developed at the laboratory. During Year 1 we found that while our systems worked well in controlled laboratory conditions, they did not work reliably enough in the relatively uncontrolled conditions of the classroom. In Year 2 we focused on the development of a robust face finder and smile detector that could provide social robots with reliable estimates of this important social behavior. The system was trained with a dataset of 70,000 images collected from the Web, containing a very wide variety of imaging conditions, races, physical appearances, etc. The new smile detector has a performance level of 96.8 % on the dataset and can run in real time at standard video rates. The system is reliable enough to be used in a wide variety of applications in real-life situations and shall be one of the perceptual primitives in the new RUBI prototype.

5.3 Learning Primitives

We developed a new approach for robots to learn to discover, in an autonomous manner, the visual appearance of objects in the world they operate [15]. In particular we conducted an experiment that demonstrated how a social robot can learn on its own to detect faces. After less than 6 minutes of interaction with the world, the robot's visual system was capable of detecting the presence of people in novel images with high accuracy (over 90 % correct).

During the 6 minutes of exposure to the world, the baby robot was never told whether or not people were present in the images, or whether people were of any particular relevance at all. It discovered that the most consistent visual explanation for the cause of the observed sensory-motor contingencies was a combination of feature detectors that happened to discriminate the presence of people very well.

5.4 Hardware Development

RUBI's robot design was inspired by Hiroshi Ishiguro's RoboVie-I humanoid [21, 28]. However, we found that the RoboVie-I design was frightening to children under 4 years of age and thus we systematically changed RUBI's appearance until children found it non-threatening. Some of the modifications included shortening the body, making it more plump, including facial expressions, clothes, a touch-screen and hair. The current RUBI prototype is a three-foot tall, pleasantly plump robot with a head, two arms, and a touch screen (See Figure 5). The exterior appearance of RUBI has been quite successful. In general the children found it non-threatening and by the end of the 13th session they exhibited a variety of social behaviors towards her including pointing, hugging, imitation, and social referencing. In Year 2 we completely redesigned RUBI's hardware while keeping her external appearance relatively unchanged.

The latest version of RUBI is constructed from two Apple PowerMac desktop computers. Each unit has two 2.5 GHz dual-core PowerPC 970MP processors with 8GB of error correcting (ECC) RAM. One machine has an 802.11g wireless card used to control the robot from a nearby wireless laptop during sessions. Currently RUBI's full software suite runs comfortably on a Mac Mini with a single 1.8 GHz Intel

Core Duo processor and 2GB of RAM, even while using two cameras with face, eye, and smile detection and color tracking running on both. Thus the two quad G5 systems give RUBI ample room to grow, and is essentially a mobile 8-node cluster. RUBI is capable of running the learning primitives mentioned in the previous section [15], which require large amounts of memory and 64-bit processing. RUBI's sensors and actuators are organized as a distributed network of control and sensing nodes, all connected to the main computer via multiple USB ports. This significantly increases communication bandwidth and avoids the bottleneck associated with having a master controller in charge of all the servos and actuators.

By far, RUBI's arms has proven the most challenging, frustrating, and elusive hardware design problem encountered in the project. The difficulty lies on the need for actuators that can handle the forces applied by children when interacting with RUBI, yet compliant and small enough to be safe. Most importantly all of this needs to be done within a very tight budget. In Year 2 we developed a streamlined 5 degrees of freedom prototype that used high end robotic RC servos. Unfortunately we ran into two problems: (1) The servos proved to be too noisy having a significant effect on the interaction with the children. (2) While they worked well in our laboratory tests, they did not survive the rigors of interaction with the children during the field studies. We are currently working on our third arm design. Critical to the new design is the issue of compliance control.

Figure 4: *A typical view from RUBI's wide angle cameras*

6. LESSONS AND CONCLUSIONS

The RUBI project was conceived as an ambitious experiment aimed at accelerating the development of robots that interact with people in everyday conditions [39, 13, 9, 21, 28]. The core principle of the project is the idea of design by immersion, i.e., the immersion of scientists, engineers and robots in the conditions of daily life as early as possible in the development process.

We believe this is important for the development of hardware, software, the discovery of the technological and scientific challenges whose solution may maximize progress, and the development of a theoretical framework for robot control. Rather than focusing on solving complex problems in the controlled conditions of laboratory environments we focused on solving simpler problems in the uncontrolled conditions of daily life.

After two years we are as convinced as ever that the design by immersion philosophy is sound and healthy. It helped us: (1) Design machine perception primitives (e.g.,

Figure 5: *RUBI teaching materials targeted by the California Results Developmental Profile from the California Department of Education.*

smile detection) that work well in field conditions, not just on the standard face datasets. (2) Develop machine learning primitives that can operate continuously in an unsupervised manner. (3) Formulate a mathematical approach to real time social interaction to handle the timing and uncertainty conditions of daily life. (4) Develop RUBIOS, a prototype software architecture for social robots. (5) Establish that long term socialization and bonding can develop between humans and robots, at least when the robot is partially controlled by a human being. (6) Identify the particularly important role of touch and haptic behaviors in the development of this bonding process. (7) Develop methods for evaluating social robot algorithms in an efficient manner in the conditions of daily life. (8) Identify the importance of studying how organisms organize behavior at multiple time scales: from reflexes to moods, emotions, and developmental processes.

We also identified lessons and limitations of the design by immersion approach as originally conceived. We were naive in the idea that we could just immerse ourselves in field conditions on a daily basis and make incremental changes until we design a "dream social robot". In practice qualitative changes are needed in the hardware and software architecture that can take months if not years away from the field. We also found that some of the intuitions initially drawn from the field sessions turned out to be misguided, perhaps setting us back in time. For example, initially we felt that self-locomotion was a critical component for progress. We invested time and effort to develop a new version of RUBI that could move autonomously about the room, only to discover that self-locomotion was perhaps distracting us away from the main focus of the project– social interaction. We under-estimated the difficulties of mechanical and sensor technology issues faced by social robots. For example, we have not managed to develop a robot arm that operates robustly in field conditions. We also underestimated the role that controlled laboratory experiments may play for testing hypotheses of interest. While the field conditions proved useful for generating hypotheses it is difficult to eliminate alternative explanations and find conclusive evidence using

field studies alone. A combination of field studies and targeted laboratory experiments may be a better strategy for progress.

Overall, we believe the RUBI project is turning out to be an exciting and useful experiment that illustrates how an immersive paradigm can help make significant progress in the emerging field of social robotics.

7. ACKNOWLEDGMENTS

8. REFERENCES

[1] http://www.sony.net/SonyInfo/QRIO.

[2] S. E. Asch. *Social Psychology*. Oxford University Press, New York, 1987 (Original work published in 1952).

[3] L. R. Bahrick and J. S. Watson. Detection of intermodal proprioceptive-visual contingency as a potential basis of self-perception in infancy. *Developmental Psychology*, 21:963–973, 1985.

[4] M. Bartlett, G. Littelwort, M. G. Frank, C. Lainscsek, I. Fasel, and J. Movellan. Recognizing facial expression: Machine learning and application to spontaneous behavior. In *IEEE International Conference on Computer Vision and Pattern Recognition*, pages 568–573, Osaka, Japan, 2005.

[5] M. Bartlett, G. Littlewort, L. C., F. I., F. M.G., and M. J.R. Fully automatic facial action recognition in spontaneous behavior. In *7th International Conference on Automatic Face and Gesture Recognition*, pages 223–228, 2006.

[6] M. S. Bartlett, G. Littlewort, C. Lainscsek, I. Fasel, and J. Movellan. Recognition of facial actions in spontaneous expressions,. *Journal of Multimedia*, in press.

[7] D. Bertsekas and S. Shreve. *Stochastic Optimal Control*. Athena Scientific, 1996.

[8] A. E. Bigelow. Infant's sensitivity to imperfect contingency in social interaction. In P. Rochat, editor, *Early social cognition: understanding others in the first months of life*, pages 241–256. LEA, New York, 1999.

[9] C. Breazeal. *Designing Sociable Robots*. MIT Press, Cambridge, MA, 2002.

[10] R. A. Brooks, C. Breazeal, M. Marjanovic, B. Scassellati, and M. M. Williamson. The Cog Project: Building a Humanoid Robot. *Lecture Notes in Artificial Intelligence*, 1562:52–87, 1999.

[11] N. J. Butko, I. Fasel, and J. R. Movellan. Learning about humans during the first 6 minutes of life. In *International Conference on Development and Learning*, Indiana, 2006.

[12] R. Cole, S. van Vuure, B. Pellom, K. Hacioglu, J. Ma, J. Movellan, S. Schwartz, D. Wade-Stein, W. Ward, and J. Yang. Perceptive animated interfaces: First steps toward a new paradigm for human computer interaction. *Proceedings of the IEEE, Special Issue on Human Computer Interaction*, 91(9):1391–1405, 2003.

[13] K. Dautenhahn. Design issues on interactive environments for children with autism. In *Proceedings International Conference on Disability, Virtual Reality and Associated Technologies (ICDVRAT)*, pages 153–161, Alghero, Sardinia, Italy, Sept. 2000.

[14] F. Tanaka., A. Cicourel., and J. R. Movellan. Socialization between toddlers and robots at an early childhood education center. under review.

[15] I. R. Fasel. *Learning to Detect Objects in Real-Time: Probabilistic Generative Approaches*. PhD thesis, UCSD, June 2006.

[16] I. Fenwick and M. D. Rice. Reliability of continuous measurement copy-testing methods. *Journal of Advertising Research*, 1991.

[17] T. Fong, I. Nourbakhsh, and K. Dautenhahn. A Survey of Socially Interactive Robots. *Robotics and Autonomous Systems*, 42(3-4):143–166, 2003.

[18] B. Fortenberry, J. Chenu, and J. R. Movellan. Rubi: A robotic platform for real-time social interaction. In *Proceedings of the International Conference on Development and Learning (ICDL04)*, The Salk Institute, San Diego, October 20, 2004.

[19] T. Ishida, Y. Kuroki, and J. Yamaguchi. Development of Mechanical System for a Small Biped Entertainment Robot. In *Proceedings of the 2003 IEEE International Workshop on Robot and Human Interactive Communication*, pages 297–302, 2003.

[20] T. Kanda, T. Hirano, D. Eaton, and H. Ishiguro. Interactive Robots as Social Partners and Peer Tutors for Children: A Field Trial. *Human-Computer Interaction*, 19(1-2):61–84, 2004.

[21] T. Kanda, H. Ishiguro, T. Ono, M. Imai, and R. Nakatsu. Development and evaluation of an interactive humanoid robot "robovie.

[22] T. Kanda, N. Miralles, M. Shiomi, T. Miyashita, I. Fasel, J. R. Movellan, and H. Ishiguro. Face-to-face interactive humanoid robot. *IEEE 2004 International Conference on Robotics and Automation*, 2004.

[23] H. Kozima, C. Nakagawa, and Y. Yasuda. Interactive Robots for Communication-care: A Case-study in Autism Therapy. In *Proceedings of the 2005 IEEE International Workshop on Robot and Human Interactive Communication*, pages 341–346, 2005.

[24] T. S. Kuhn. *The Structure of Scientific Revolutions*. University of Chicago Press, Chicago, 1996.

[25] Y. Kuroki, T. Fukushima, K. Nagasaka, T. Moridaira, T. T. Doi, and J. Yamaguchi. A Small Biped Entertainment Robot Exploring Human-Robot Interactive Applications. In *Proceedings of the 2003 IEEE International Workshop on Robot and Human Interactive Communication*, pages 303–308, 2003.

[26] G. Littlewort, M. Bartlett, I. Fasel, J. Susskind, and J. R. Movellan. An automatic system for measuring facial expression in video. *Image and Vision Computing.*, in press.

[27] D. Marr. *Vision*. W.H. Freeman and Company, San Francisco, 1982.

[28] T. Minato, M. Shimada, H. Ishiguro, and S. Itakura. Development of an android robot for studying human-robot interaction.

[29] J. R. Movellan. An infomax controller for real time detection of contingency. In *Proceedings of the International Conference on Development and Learning (ICDL05)*, Osaka, Japan, 2005.

[30] J. R. Movellan, F. Tanaka, B. Fortenberry, and K. Aisaka. The RUBI project: Origins, principles and first steps. In *Proceedings of the International Conference on Development and Learning (ICDL05)*, Osaka, Japan, 2005.

[31] J. R. Movellan and J. S. Watson. The development of gaze following as a Bayesian systems identification problem. In *Proceedings of the International Conference on Development and Learning (ICDL02)*. IEEE, 2002.

[32] Y. Nagai, M. Asada, and K. Hosoda. Learning for Joint Attention Helped by Functional Development. *Advanced Robotics*, 20(10):1165–1181, 2006.

[33] A. Pentland. Socially Aware Computation and Communication. *IEEE Computer*, 38(3):33–40, 2005.

[34] J. Peter H. Kahn, B. Friedman, D. R. Perez-Granados, and N. G. Freier. Robotic Pets in the Lives of Preschool Children. *Interaction Studies*, 7(3):405–436, 2006.

[35] R. W. Picard. *Affective Computing*. The MIT Press, 1997.

[36] M. Polsfuss and M. Hess. Liking through moment-to-moment evaluation: Identifying key selling segments in advertisement. *Advances in consumer research*, 18:540–544, 1991.

[37] M. J. R. and J. S. Watson. Perception of directional attention. In *Infant Behavior and Development: Abstracts of the 6th International Conference on Infant Studies*, NJ, 1987. Ablex.

[38] 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*, 4(2):105–120, 2005.

[39] B. Scassellati. *Foundations for a Theory of Mind for a Humanoid Robot*. PhD thesis, MIT Department of Computer Science and Electrical Engineering, 2001.

[40] J. Susskind, J. Hershey, and J. Movellan. Exact inference in robots using topographical uncertainty maps. In *Proceedings of the second international conference on development and learning (ICDL04)*, The Salk Institute, San Diego, October 20, 2004.

[41] J. Susskind, G. Littlewort, M. Bartlett, J. Movellan, and A. Anderson. Human and computer recognition of facial expressions of emotion. *Neuropsychologia.*, in press.

[42] 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: (Best Paper Award). In *Proceedings of the IEEE International Workshop on Robot and Human Interactive Communication (RO-MAN)*, pages 223–228, Nashville, USA, 2005.

[43] F. Tanaka, B. Fortenberry, K. Aisaka, and J. R. Movellan. Plans for developing real-time dance interaction between QRIO and toddlers in a classroom environment. In *Proceedings of the International Conference on Development and Learning (ICDL05)*, Osaka, Japan, 2005.

[44] F. Tanaka, J. R. Movellan, B. Fortenberry, and K. Aisaka. Daily HRI evaluation at a classroom environment: Reports from dance interaction experiments. In *Proceedings of the 2006 Conference on Human-Robot Interaction (HRI 06)*, Salt Lake City, 2006.

[45] F. Tanaka and H. Suzuki. Dance Interaction with QRIO: A Case Study for Non-boring Interaction by Using an Entrainment Ensemble Model. In *Proceedings of the 2004 IEEE International Workshop on Robot and Human Interactive Communication*, pages 419–424, 2004.

[46] J. S. Watson. The perception of contingency as a determinant of social responsiveness. In E. B. Thoman, editor, *Origins of the Infant's Social Responsiveness*, pages 33–64. LEA, New York, 1979.

[47] J. S. Watson. Contingency perception in early social development. In F. T. M. and F. N. A., editors, *Social perception in infants*, pages 157–176. Ables, New Jersey, 1985.

[48] J. Weng, J. McClelland, A. Pentland, O. Sporns, I. Stockman, M. Sur, and E. Thelen. Autonomous Mental Development by Robots and Animals. *Science*, 291(5504):599–600, 2000.

[49] D. Wolpert, K. Doya, and M. Kawato. A unifying computational framework for motor control and social interaction. In F. C and W. DM, editors, *The Neuroscience of Social Interaction*, pages 305–322. Oxford University Press, Oxford, UK, 2004.