IEEE CIS

AMD Newsletter

Robotics
Psychology
Neuroscience
Machine Intelligence

The Newsletter of the Autonomous Mental Development Technical Committee

VOLUME 9, NUMBER 1 ISSN 1550-1914 April 2012

Editorial

New fundamental concepts, experimental methods and techniques have been growing recently about the understanding of developmental social learning in humans and robots. They are at the crossroads of research in developmental robotics, human-robot interaction, and the psychology and neuroscience of human social learning. This newsletter features a stimulating dialog illustrating the high-importance of these concepts and techniques both for building machines capable of learning in interaction with non-engineer humans across an extended time-scale, but also for better understanding human development. Together with an initiation and a synthesis by Katharina Rohlfing and Britta Wrede, the newsletter features contributions from Yukie Nagaï, Andrea Thomaz, Lakshmi Gogate, Adriana Tapus, Stevo Bozinoski, Anthony Morse, Thomas Hannagan, Rachel Wu, and Helge Ritter.

This month, a new dialog initiation by Denis Mareschal challenges existing computational models of development by outlining three processes which he argues are fundamental in child development but mostly absent from models:

- Epigenetic dynamics where gene expression is modified during ontogeny in interaction with the environment;
- Body growth and the non-optimality of sensor fusion in infants;
- The emotional mechanism associated with social learning, which control brain plasticity and help infants to select who to learn from.

Those of you interested in reacting to this dialog initiation are welcome to submit a response (contact pierre-yves.oudeyer@inria.fr) by September 1st, 2012. The length of each response must be between 300 and 500 words (including references).

— Pierre-Yves Oudeyer, INRIA, Editor

Welcome Message from the New Chair of AMD Technical Committee

Dear colleagues, it is an honour, but also a great responsibility, to have to follow on the steps of Minoru Asada and the previous chairs, and lead the AMD Technical Committee for 2012. To allow continuity with previous leadership, I have asked Matt Schlesinger to remain Deputy-Chair for the Americas region, and asked Yukie Nagai to stand in as Deputy Chair for the Asia region. And I would like to express my warmest thanks to Minoru Asada and Jochen Triesch for their invaluable work in the past two years.

This 2012 promises to be a very busy year for our community. Below are the main activities organised by AMD staff and planned for the coming months.

- ICIS 2012 Preconference Workshop on Developmental Robotics (6 June 2012), in Minnesota USA. This is organised in conjunction with ICIS-2012, the International Conference on Infant Studies, and will provide a showcase of developmental robotics work to developmental psychologists. The workshop is organised by Matt Schlesinger and Jochen Triesch.
- Brain-Mind Institute’s Summer School (June-August 2012) and the ICBM, International Conference on Brain-Mind (14-15 July 2012) in East Lansing, Michigan USA. These are the result of John Weng’s enthusiastic creation of the Brain-Mind Institute.
- WCCI-2012 Special Session on Bio-Inspired Developmental Mechanisms (10-15 June 2012), in Brisbane Australia.

- The most important event for the AMD community this year will be the 2012 IEEE ICDL-EpiRob Conference, to be held in San Diego, California, on 7-9 November 2012. This is chaired by Javier Movellan, with the support of Matt Schlesinger and Jochen Triesch as co-general chairs. The deadline for papers submission is 15 June 2012.

Here is an update on the IEEE AMD Technical Committee. My first aim as Chair has been to extend the membership of the AMD Technical Committee, and in particular to balance the geographical representation and reach of the AMD members. We have increased the AMD TC members from 65 to 94, and we are lucky now to have members from countries such as Australia, Brazil, Greece, India, Russia and Turkey. A full list of members is available in the AMDTC website www.ieee-cis.org/technical/amdtc/. We now also have a mailing list ieee-amdtc@googlegroups.com

The second aim of the AMD TC is to look at how we can best serve our community through the Task Forces. In the coming months we will be working on a full review of the TC sub-group organisation, e.g. to consider adding new Task Forces on Human-Robot Interaction, Neuroscience, Web presence, and the revising and regrouping of existing task forces. We need volunteers to lead and contribute to current and new Task Forces, so please feel free to offer your help.

And finally, remember to support our journal, the IEEE Transactions for Autonomous Mental Development, with the submission of your best papers (no need to try first Science and Nature ;-), and the reading (and citation!) of the work published in the journal. IEEE TAMD has now been selected for inclusion in the main impact factor and citation databases.

I look forward to working with you all in the coming year. And I welcome and strongly encourage suggestions for new activities and initiatives for the AMD community at large, as well as for the AMD Technical Committee. Feel free to drop an email with new ideas to acangelosi@plymouth.ac.uk.

— Angelo Cangelosi, the new chair of AMD TC

Message from the Former Chair of AMD Technical Committee

I would like to thank all the continuous contributions by members of the AMD community to both the activities of our AMD TC and the growth of our scientific domain. I have finished my term as AMD TC chair for the years 2010 and 2011. During these two years, we had several major events which were highly successful. The biggest one, in 2011, was the joint organization of ICDL and EpiRob conferences, in Frankfurt, Germany. I would like to thank the general chairs of the conference, Jochen Triesch (the former vice chair of AMD TC) and Angelo Cangelosi (the current chair of the AMD TC). This year, we will have the second one at UCSD, hosted by Javier Movellan, who will be general chair together with Matthew Schlesinger and Jochen Triesch. We will have the third one in Osaka, Japan, expecting it as a symbol of the restoration from the quake, and Yukie Nagai shall work for it. Another important target of our activities is the organization of workshops (and possibly an ICDL-Epirob conference in the future) co-located with developmental and neuroscience conferences. As the first step, Matthew Schlesinger (the former vice chair of AMD TC) is organizing ICIS (International Conference on Infant Study) 2012 Preconference Workshop on Developmental Robotics.

As mentioned, Angelo Cangelosi is the new chair of the AMD TC, and he has already started his work in order to further improve the quality and impact of our activities, through for example an update of the task forces. He and his team are very active, promising, and expected to address challenges ahead, and I encourage volunteers to contribute and help him in this endeavour. Thanks again for your contributions to the past, current, and future activities of the AMD TC.

— Minoru Asada, the former chair of AMD TC

Dialog Column

What Novel Scientific and Technological Questions Does Developmental Robotics Bring to HRI? Are we Ready for a Loop?

Katharina J. Rohlfing & Britta Wrede
Center of Excellence Cognitive Interaction Technology (CITEC) Bielefeld University, Germany

The notion of embodied systems that enable multi-modal interaction and thus facilitate learning grounded in sensorimotor experience is a current overarching paradigm in robotics. However, we argue that taking a developmental stance should take us beyond active but self-directed cognition [7]; developmental robotics should provide us with insights into the advantages of social learning. Thus, it should be questioned how systems can take advantage of the input and incremental interaction with the social and physical environment. Only then we will find new dimensions enabling learning processes on robots that help to overcome current limitations of generalization. However, we suspect that our current methods – accessing children's development as well as designing learning systems – hinder us from asking the right questions.

In human developmental research, the view of active cognition [7] is supplemented by a view focusing on specifically designed input and the fact that children's minds allow taking advantage of it. Research [2,10,17,23] persuasively shows that cognition is not only self-directed but also distributed – over a system of people and objects within an environment ([17], p. 97). In this system, crucial learning information is ostensive [3] and reduced; we argue that robotic systems that are sensitive towards such input benefit from this specifically placed reduced information. We have implemented a first approach towards such a tutor spotter suggesting that it can induce tutoring behavior by users [13].

However, research on human infants has shown that effective learning requires more than just providing social input. The social information has to be flexible as it should be co-constructed online ([4]) and contingent with the child's feedback (e.g. [14]). Our own research contributes to the argument that it is not only the in- put but rather the interplay between the input and the feedback of the participants that enables the learner to take advantage of tutoring: in [15, 19], we showed that learners provide feedback in the form of e.g. their eye-gaze (signaling their attention or anticipating subsequent actions) which shapes the way input is provided (e.g.[18], [15]) and it is crucial for robotic systems to elicit multimodal tutoring input [6]. This, however, is only the first step towards understanding the power of interaction and how, within this exchange, the specifically tailored input influences what is learned.

Based on research results pointing to the loop, we question the capability of current robotics approaches to learn and generalize actions as well as language in embodied systems. Specifically we see shortcomings with respect to the following:

1. Current representations consider knowledge as a static entity, where incrementality is interpreted as adding new data points (e.g. [16]) or (sometimes) new classes (e.g. [9,1]). However, when we take more seriously the fact that learning takes place in an interaction loop [15,20], we have to design representations that are inherently dynamic and store knowledge not as a binary but a dynamic state to which the environment and experiences contribute. First solutions for emerging hierarchical structures have been presented by e.g. Tani and colleagues [21] but in these approaches, structure emerges at different levels of a hierarchy with higher levels serving as sequencing concepts over lower level motor primitives rather than from vague to concrete.

2. Learning systems support short-term memory. However, learning systems need to be equipped with long-term memory that facilitates knowledge assimilation and consolidation processes (cf. [8] for a memory-based language learning approach).

3. Current systems are restrictive with respect to the pragmatic frame of interaction. For example, either the tutor's input or self-exploration are the sources of learning. However, a learning system needs to be able to switch roles, thus becoming a tutor as

Dialog Column

well as mixing different sources of its learning in order to allow for self-reflexive processes consolidating accumulated knowledge further.

4. Supervised learning algorithms are not able to take qualitative feedback of a tutor into account. However, we need to consider feedback that goes beyond binary reinforcement signals but rather emphasizes specific parts of the learner‘s behavior, such as the manner in which an object is grasped or the goal where it has to be put.

The complementary, methodical monopoly in developmental research let the studies focus only on one side of the learning process (either the learner or the tutor) and falls short of accessing the interplay between the tutor‘s behavior and the learner‘s feedback. Exceptions to the monopoly are studies that manage to encompass the loop in interaction in a systematic way ([11-12]) and thus provide a comprehensive view on children‘s learning. We think that a bridge between qualitative and quantitative methods of analysis offers a solution to advance current approaches: Qualitative methods like Conversational Analysis can aid us in identifying the means that the participants use to signal feedback [19]; quantitative methods give us evidence about the significance of different types of feedback on the shape of the interaction loop [12]. Using multiple and new methods ([22]) will foster insights into how inter- and intra-personal coordination are related to each other [5] as the interaction unfolds and how it drives long-term learning processes.

References:

[1] Allyon, I. C., Heckmann, M, Denecke, A., Wrede, B. & Goerick, C. (2010) Incremental word learning using large margin discriminative training and variance floor estimation. In: Proceedings of the INTERSPEECH.
[2] Brand, R. J. & Tapscott, S. (2007) Acoustic packaging of action sequences by infants. In: Infancy 12: 321–332.
Beer, Randall D. (2000) Dynamical approaches to cognitive science. In: Trends in Cognitive Science 4: 91–99.
[3] Csibra, G. & Gergely, G. (2009) Natural pedagogy. In: Trends in Cognitive Science 13: 148–153.
[4] Colins, S. & Marková, I. (1999) Interaction between impaired and unimpaired speakers: Inter-subjectivity and the interplay of culturally shared and situation specific knowledge. In: British Journal of Social Psychology 38: 339-368.
[5] De Jaegher, H., Di Paolo, E., & Gallagher, S. (2010) Can social interaction constitute social cognition? In: Trends in Cognitive Sciences 14: 441– 447.
[6] Fischer, K., Foth, K., Rohlfing, K. J., & Wrede, B. (2011) Mindul tutors: Linguistic choice and action demonstration in speech to infants and robots. In: Interaction Studies 12: 134–161.
Frith, C. D. & Frith, U. (2011) Mechanisms of social cognition. In: Annual Review of Psychology 63: 8.1–8.27.
[7] Gelman, S. A. (2009) Learning from others: Children's construction of concepts. In: Annual Review of Psychology 60: 115–140.
[8] Gläser, C. & Joublin, F. (2010) A computational model for grounding words in the perception of agents. In: Proceedings of the IEEE International Conference on Development and Learning 26–32.
[9] Gnadt, W. & Grossberg, S. (2008) SOVEREIGN: An autonomous neural system for incrementally learning planned action sequences to navigate towards a rewarded goal. In: Neural Networks 21: 699–758.
[10] Gogate, L., Bahrick, L. & Watson, J. (2000) A study of multimodal motherese: The role of temporal synchrony between verbal labels and gestures. In: Child Development 71: 878–894.
[11] Hofer, T., Hohenberger, A., Hauf, P. & Aschersleben, G. (2008) The link between maternal interaction style and infant action understanding. In: Infant Behavior and Development 31: 115–126.
[12] Koterba, E. A. & Iverson, J. M. (2009) Investigating motionese: The effect of infant-directed action on infants' attention and object exploration. In: Infant Behavior and Development 32: 437–444.
[13] Lohan, K. S., Pitsch, K., Rohlfing, K. J., Fischer, K., Saunders, J., Lehmann, H., Nehaniv, C. & Wrede, B. (2011) Contingency allows the robot to spot the tutor and to learn from interaction. In: Proceedings of the IEEE International Conference on Development and Learning and Epigentic Robotcs (ICDL-EpiRob).
[14] Markova, G. & Legerstee, M. (2006) Contingency, imitation, and affect sharing: Foundations of infants' social awareness. In: Developmental Psy- chology 42: 132–141.
[15] Pitsch, K., Vollmer, A.-L., Fritsch, J., Wrede, B., Rohlfing, K. J. & Sagerer, G. (2009) On the loop of action modification and the recipient’s gaze in adult-child interaction. In: Proceedings of the International Conference on Speech and Gesture in Interaction, September (GESPIN).

Dialog Column

[16] Rodrigues, C., Gérard, P. & Rouveirol, C. (2010) Incremental learning of relational action models in noisy environments. In: Proceedings of the Int.Conf. on Inductive Logic Programming (ILP 2010).
[17] Rohlfing, K. J., Fritsch, J., Wrede, B. & Jungmann, T. (2006) How can multimodal cues from child-directed interaction reduce learning complexity in robots? Advanced Robotics 20: 1183–1199.
[18] Smith, N. A. & Trainor, L. J. (2008) Infant-directed speech is modulated by infant feedback. In: Infancy 13: 410–420.
[19] Vollmer, A.-L., Pitsch, K., Lohan, K., Fritsch, J., Rohlfing, K. J. & Wrede, B. (2010) Developing feedback: How children of different age contribute to an interaction with adults. In: Proceedings of the IEEE 9th International Conference on Development and Learning: 76–81.
[20] Wrede, B., Kopp, S., Rohlfing, K. J., Lohse, M. & Muhl, C. (2010) Appropriate feedback in asymmetric interactions. In: Journal of Pragmatics 42:2369–2384.
[21] Yamashita, Y. & Tani, J. (2008) Emergence of functional hierarchy in a multiple timescale neural network model: a humanoid robot experiment, In:PLoS Computational Biology 11.
[22] Yu, C., Smith, L.B., Shen, H., Pereira, A.F., & Smith, T.G. (2009) Active information selection: Visual attention through the hands. In: IEEE Transactions on Autonomous Mental Development 1: 141–151.
[23] Zukow-Goldring, P. (1996) Sensitive caregiving fosters the comprehension of speech: When gestures speak lauder than words. In: Early Development and Parenting 5: 195–211.

Mutual Shaping between Tutor’s Scaffolding and Robot’s Development

Yukie Nagai
Graduate School of Engineering, Osaka University, Japan

Rohlfing and Wrede point out throughout their dialog initiation that there is still a gap between robot learning and social interaction: current robots are not endowed with the ability to take advantage of social learning and thus not yet ready for an interaction loop. Although many studies have investigated robot learning with human scaffolding (e.g., [1,3-5]), they have focused either on robot learners (e.g., how robots learn from tutors) or on human tutors (e.g., how tutors want to teach robots).

I agree that a missing link is the interplay between a tutor’s input and a robot’s feedback. A tutor should adapt the complexity of input so as to facilitate the robot’s learning (i.e., tutors’ scaffolding) while the robot should provide appropriate feedback to elicit proper scaffolding by the tutor (i.e., robots’ development).

Toward modeling such mutual shaping between scaffolding and development, we recently started a close analysis of caregiver-infant interaction [2]. Inspired by [6], we have collected motion data of caregivers and infants during their dynamic interaction, and measured the information flow (i.e., how much a motion influences another motion) between and within a caregiver and an infant. For example, an infant’s development (e.g., improvement of body coordination) can be detected as an increase in the information flow within the infant. A caregiver’s behavioral adaptation to an infant’s development can be detected as a change in his/her within-information flow (e.g., the higher the flow is, the more highly coordinated the caregiver’s motions are). Furthermore, social contingency can be measured as the information flow between a caregiver and an infant. Our latest results reveal the mechanism of mutual shaping between caregivers’ scaffolding and infants’ development, and show a further potential to uncover fine-grained social learning.

As Rohlfing and Wrede mention, there are still several issues to address in order to enable robots to take advantage of social learning. The biggest remaining challenge would be to integrate different aspects of social learning such as the design of learning systems and the coordination of interaction. I hope our interdisciplinary research in developmental robotics would break though the current limitations.

Dialog Column

References:

[1] Nagai, Y., Asada, M., and Hosoda, K. (2006) Learning for joint attention helped by functional development. Advanced Robotics, vol. 20, no. 10, pp. 1165-1181.
[2] Nagai, Y. (2012) Reading Intentions from Motionese: Analyzing and Designing Caregiver-Infant Interaction. Invited talk at the Workshop on "Reading intentions: From children to robots," Lund, Sweden.
[3] Calinon, S. and Billard, A. (2008) A framework integrating statistical and social cues to teach a humanoid robot new skills. In Proceedings of the ICRA 2008 Workshop on Social Interaction with Intelligent Indoor Robots.
[4] Thomaz, A. L. and Breazeal, C. (2008) Teachable robots: Understanding human teaching behavior to build more effective robot learners. Artificial Intelligence Journal, vol. 172, pp. 716-737.
[5] Saunders, J., Nehaniv, C. L., Dautenhahn, K., and Alissandrakis A. (2007) Self-imitation and Environmental Scaffolding for Robot Teaching. International Journal of Advanced Robotic Systems, vol. 4, no. 1, pp. 109-124.
[6] Pereira, A. F., Smith, L. B., and Yu, C. (2008) Social coordination in toddler's word learning: interacting systems of perception and action. Connection Science, vol. 20, no. 2–3, pp. 73–89.

Socially Guided Robot Learning

Andrea L. Thomaz
Georgia Institute of Technology, USA

While state of the art robot learning technology is not ready to participate in the kind of open-ended social learning interactions that children do, I think we are in a position to make important first steps in this direction. This is the goal of my own research in *Socially Guided Machine Learning*, which aims to computationally model mechanisms of human social learning to achieve robots that are intuitive for people to teach.

Much of our work has focused on point (4), that supervised learning techniques are not designed for input from naive humans. We often begin with an investigation into the feasibility of a particular machine learning interaction, which leads to a series of research questions around re-designing both the interaction and the algorithm to better suit learning with end-users. The following are three examples:

Reinforcement Learning (RL)
A common approach for incorporating human input into RL is reward shaping, letting the human directly control the reward signal to the agent [1,2]. We began with this type of interaction and iteratively designed new interfaces and algorithms, through experiments with humans, to support the ways that we observed people trying to teach. For example, in addition to reward, people want to guide the agent's attention during learning as a lightweight form of action advice. And if the agent infers that negative feedback implies an "undo" request, this leads to a 50% speedup in learning.

Robot Learning from Demonstration (LfD)
Much of the field of LfD is motivated by the notion of robots learning from end-users. A survey of LfD [4] shows a range of different input schemes with very different interactions for the end-user (e.g., teleoperating, motion capture, or moving a robot kinesthetically to provide learning data). However, the field lacks an understanding of the usability of various input mechanisms. Our aim is to create LfD techniques that end-users find natural and intuitive. We have focused on one popular input mode, kinesthetic teaching. Recently we have shown the benefits of allowing end users to demonstrate full trajectories versus sparse "keyframe" trajectories, finding each is useful for different kinds of skills [5]. And since existing LfD algorithms are designed only for full trajectories, we have introduced a new framework for keyframe-based LfD [6].

Dialog Column

Active Learning for robots

Active Learning (AL) is a machine learning paradigm in which the learner queries the supervisor, requesting particular learning input. Thus, an obvious candidate for social robot learning. Our initial investigation compared an AL querying mechanism to a passive supervised learning method (with a human teacher as the supervisor). We find that active learning has the potential to greatly improve the interaction from a learning performance perspective, but results suggest it may have an inappropriate balance of control [7]. One finding from this work was that people's compliance to the robot's queries was dependent on their believing it was a "good" question. Thus, we have adapted three types of questions from the AL literature (feature, label, and demo queries) to make them appropriate for embodied interactive learning, and verified their utility with human partners [8].

References:

[1] Ng, A. Y., Harada, D., Russell, S. J. Policy Invariance Under Reward Transformations: Theory and Application to Reward Shaping, Proceedings of the Sixteenth International Conference on Machine Learning, p.278-287, June 27-30, 1999

[2] Isbell, C., Shelton, C., Kearns, M., Singh, S. and Stone, P. Cobot: A social reinforcement learning agent. 5th International Conference on Autonomous Agents, 2001.

[3] Thomaz, A.L. and Breazeal, C. Teachable robots: Understanding human teaching behavior to build more effective robot learners. Artificial Intelligence Journal, 172:716-737, 2008.

[4] Argall, S. Chernova, S., Veloso, M. and Browning, B. A Survey of Robot Learning from Demonstration. Robotics and Autonomous Systems. Vol. 57, No. 5, pages 469-483, 2009.

[5] Akgun, B., Cakmak, M., Yoo, J.W. and Thomaz, A.L.. Trajectories and Keyframes for Kinesthetic Teaching: A Human-Robot Interaction Perspective. International Conference on Human-Robot Interaction, 2012.

[6] Akgun, B., Cakmak, M., Jiang, K. and Thomaz, A.L. Keyframe-based Learning from Demonstration. International Journal on Social Robotics, to appear in 2012.

[7] Cakmak, M., Chao, C. and Thomaz, A.L.. Designing Interactions for Robot Active Learners. IEEE Transactions on Autonomous Mental Development, March, 2010.

[8] Cakmak, M. and Thomaz, A.L.. Designing Robot Learners that Ask Good Questions. International Conference on Human-Robot Interaction, 2012.

The Relative Novelty of Embodied Systems in Human and Human-Robot Interaction

Lakshmi J. Gogate
Florida Gulf Coast University, Fort Myers, USA.

Rohlfing and Wrede aptly take a developmental stance to consider an important set of issues pertaining to learning within artificial and natural embodied systems. Their main argument is that embodied systems can inherently take advantage of the input and incremental interaction with the social and physical environment. Thus, natural and artificial systems can benefit from not just the input in isolation but "the [reciprocal] interplay between the input and the feedback of the participants [learners]". Also emphasized in their initiated dialog is that the learner can benefit from reduced input that is multisensory. A final issue underscored is that incremental learning by the participant is possible with incremental increases in complexity of the input and incremental readiness of the participant to receive this input.

It must be acknowledged, however, that the main argument and several ensuing questions are not entirely novel: They have been extensively discussed in the developmental literature and have been empirically tested in developmental psychology and in developmental robotics research. The present dialogue therefore could benefit greatly from further in-depth consideration of three related areas of research. First, it would be useful to consider the extant developmental theory on embodied and extended cognition where intelligence is distributed across brain, body and environment [1, 2, 3, 4]. Second, let us consider the empirical evidence from language development research. These studies show that the interplay between the tutor's behavior and the

Dialog Column

learner's feedback, specifically infant's gaze-switching behavior from mother to object during her naming in temporal synchrony with shaking and looming object motions contributes to word learning [5, 6, 7, 8]. These studies combine observational and experimental techniques to further elucidate the interplay between tutor and learner. Gogate and Hollich [3] provide several other examples of this interplay between tutor and learner during language learning [e.g., 9, 10]. Finally, let us consider some studies of human-robot interaction which eloquently demonstrate the interplay between tutor and learner. As a case in point, work by Kozima and colleagues shows that, children with autism benefit socially from contingent interactions with a robot [11]. Similarly, work by Poulin-Dubois and colleagues illustrates that typically developing infants can understand intentionality better from a robot "caregiver" if its labeling is contingent with its actions [12]. In addition, Nagai and Rohlfing have demonstrated that a saliency-based attention model shows greater attention to a parent's infant-directed talking face during real-time parent–infant interaction with objects than to the same parent's face during parent–adult interaction [13].

References:

[1] Clark, A. (1997) *Being there: Putting brain, body, and world together again.* Cambridge, MA: MIT Press.

[2] Clark, A., & Chalmers, D. (1998) The extended mind. *Analysis, 58,* 7–19.

[3] Gogate, L.J. & Hollich, G.J. (2010) Invariance detection within an interactive system: A perceptual gateway to language development. *Psychological Review*, 171(2), 496-516.

[4] Pfiefer, R., & Scheier, C. (1999) *Understanding intelligence.* Cambridge, MA: MIT Press.

[5] Gogate, L. J., Bolzani, L. H., & Betancourt, E. (2006) Attention to maternal multimodal naming by 6- to 8-month-old infants and learning of word-object relations. *Infancy, 9,* 259–288.

[6] Gogate, L. J. (2010) Learning of syllable-object relations by preverbal infants: The role of temporal synchrony and syllable distinctiveness. *Journal of Experimental Child Psychology, 103,* 178–197.

[7] Matatyaho, D., & Gogate, L. J. (2008) Type of maternal object motion during synchronous naming predicts preverbal infants' learning of word-object relations. *Infancy, 13,* 172–184.

[8] Matatyaho, D., Mason, Z., & Gogate, L. J. (2007) Word learning by 8-month-old infants: The role of object motion and synchrony. *Proceedings of the 7th International Conference on Epigenetic Robotics: Lund University Cognitive Studies, 135,* 201–202.

[9] Cameron-Faulkner, T., Lieven, E., & Tomasello, M. (2003) A construction based analysis of child directed speech, *Cognitive Science, 27,* 843–873.

[10] Adamson, L., & Bakeman, R. (2006) The development of displaced speech in early mother-child conversations. *Child Development, 77,* 186–200.

[11] Kozima, H., Nakagawa, C., & Yasuda, Y. (2005) Interactive robots for communication-care: A case-study in autism therapy. *IEEE International Workshop on Robot and Human Interactive Communication,* (ROMAN-2005), 341–346.

[12] Poulin-Dubois, D. (2006) Infants' concept of mentalistic agent: Who can have intentions and goals? In S. Itakura, K. Hiraki, A. Arita, K. Kamewari, R. Baillargeon, & D. Poulin-Dubois, *Infants' Conceptual Representation of Intentional Agent.* Symposium conducted at the International Conference on Infant Studies, Kyoto, Japan.

[13] Nagai, Y., & Rohlfing, K. J. (2007) *Can motionese tell infants and robots "what to imitate"?* Meeting of the Society for Artificial Intelligence and the Simulation of Behaviour, Bristol, England.

Dialog Column

Is Developmental Robotics a Solution for Socially Assistive Robotics?

Adriana Tapus
Cognitive Robotics Lab ENSTA-ParisTech, Paris, France

The Human-Robot Interaction (HRI) field for assistive applications focuses on how to provide long-term/lifelong social interaction for vulnerable populations (e.g., children with autism, post-stroke patients, and individuals suffering from cognitive impairments) [1]. Due to the sensitive nature of their interactions with humans, socially assistive robots need a thorough and targeted training, similar to that received by nurses/trainers/therapists, before they are released into the wild. They also need to continuously and incrementally learn and adapt their behavior to the user's profile (i.e., personality, preferences, and disability) and to the environmental changes in order to deliver a personalized, engaging, and motivating social interaction and useful feedback to their "users".

Different perceptual modalities and cognitive and behavioral capabilities need to be explored so that a robot can develop through continuous interaction with, learning of, and adaptation in the social environment. The robot's decisional abilities require taking into account context, user's profile and disability level while performing its tasks [1]. It has to take initiative to establish and conduct a fruitful therapeutical session with humans, and change its interaction styles depending on context and scenarios [2,3]. The point raised by Rohlfing and Wrede regarding the necessity of a long-term memory to facilitate knowledge assimilation and consolidation is one that I also share. I would go even further and argue that sharing learned concepts between robots, in a distributed fashion, and make some of that knowledge innate to the next generation is extremely valuable. However, due to the sensitive nature of the tasks and interactions between socially assistive robots and the vulnerable users that require their assistance, I believe that a traditional machine learning approach that allows for manual intervention in the learning process, as opposed to the "skull closed" approach advocated by the autonomous mental development community, is more appropriate.

The multimodal sensing capabilities that are required to move and act for long periods of time in continuously changing, human-centered environments have highlighted the importance of the autonomous mental development of robots. Moreover, developing robots capable of expressing intentionality and spontaneity in social interaction is another problem, where the developmental process is more appropriate [4]. Nevertheless, in the assistive context, one difficulty encountered in lifelong learning robots is to measure what the robot has learned and its evolution in time. A qualitative and quantitative answer to this question could provide the possibility for the human to act on and to influence the robot decisional processes and vice versa.

Rohlfing and Wrede discuss also the importance of incorporating some high-level abstract representation of tasks and object affordances in the system. This implies having or constructing some semantic representation and/or ontological schema. Most of the existing works in robotics have tried to learn the affordances [5] and very little have addressed the key problem of inferring affordance parameters from multimodal perceptual measurements [6]. This is of high importance and more research should be pursued in this direction.

To conclude this short response, recent works in socially assistive robotics show the importance of long-term interaction in having a personalized adaptive behavior based on the experience achieved and the interaction episodes [7,3]. In this context, developmental robotics, still in its infancy, can perhaps bring some new solutions to assistive robotics and therefore an interdisciplinary collaboration that creates the marriage between these two fields is required.

References:

[1] Tapus, A., Tapus, C. and Mataric, M. J. (2008) User-robot personality matching and robot behavior adaptation for post-stroke rehabilitation therapy, *Intelligent Service Robotics,* vol. 1, no. 2, pp. 169–183.
[2] Tapus, A. and Mataric, M. J. (2008) Socially assistive robotic music therapist for maintaining attention of older adults with cognitive impairments, in Proceedings *of AAAI Fall Symposium AI in Eldercare: New Solutions to Old Problem,* Washington D.C., USA.

Dialog Column

[3] Tapus, A., Tapus, C. and Mataric, M. J. (2010) Long term learning and online robot behavior adaptation for individual swith physical and cognitive impairments, in *In Field and Service Robotics, Results of the 7th International Conference,* A. Kelly, K. Iagnemma, and A. Howard, Eds. Series: Springer Tracts in Advanced Robotics, vol. 62, pp. 389–401.
[4] Baranes, A. and Oudeyer, P.-Y. (2009) R-IAC: Robust intrinsically motivated exploration and active learning, *IEEE Transactions on Autonomous Mental Development,* vol. 1, no. 3, pp. 155–169.
[5] Rome, E., Hertzberg, J., Dorffner, G. and Doherty, P. (2006) Towards affordance-based robot control, in ser. Dagstuhl Seminar Proceedings, E. Rome, P. Doherty, G. Dorffner, and J. Hertzberg, Eds., no. 06231. Dagstuhl, Germany: Internationales Begegnungs- und Forschungszentrum fur Informatik (IBFI), Dagstuhl, Germany.
[6] Hermans, T., Rehg, J. M. and Bobick, A. (2011) Affordance prediction via learned object attributes, in IEEE International Conference on Robotics and Automation (ICRA): Workshop on Semantic Perception, Mapping, and Exploration, Shanghai, China.
[7] Demiris, Y. (2009) Knowing when to assist: Developmental issues in lifelong assistive robotics, in Proceedings of the Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC), Minneapolis, MN, pp. 3357–3360.

Bring on The Loop

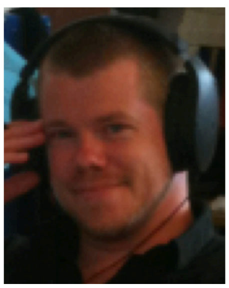

Anthony Morse, Cognition Institute, Centre for Robotics and Neural Systems, Plymouth University, UK

Thomas Hannagan, Laboratoire de Sciences Cognitives et Psycholinguistique, EHESS/CNRS/DEC-ENS, France

Rachel Wu, Centre for Brain and Cognitive Development, Birkbeck, University of London, UK

We concur with Rohlfing and Wrede in that developmental robotics should both inform, and be informed by the fields of social and infant cognition. Embodied cognition has, to date, focused on agent-environment interaction to the exclusion of social interactions, and so has missed this incredibly rich (and often simplifying) source of information. We would, however, like to make some qualifications to the picture of infant cognition described, and the questions raised. We do not contest the role of online social feedback and contingent interactions in infants learning (perhaps most compellingly demonstrated when it is lacking, altered, or misinterpreted, as in autism spectrum disorders). Recently however, a number of studies have shown that even by 8-9 months of age the mere presence of social cues in videos (e.g., a friendly face happily addressing the infant) improves the likelihood of learning audio-visual events compared to non-social cues (e.g., flashing lights; [1]). These social cues fail to meet the criteria set by Rohlfing and Wrede: they are neither flexible, co-constructed online, nor contingent with the infants behavior. However, it is likely that they would depend on previous appropriate feedback experience.

This newly emerging literature shows that beyond interactivity, social cues are useful for infants. This in turn points to different avenues of research for developmental robotics in addition to those outlined by Rohlfing and Wrede. For instance, new research questions can address how, why, and when social cues (e.g., eye gaze, motionese, infant-directed speech, imitation) can improve infants learning of both social and non-social targets (e.g. [2]). The authors also call for a cognitive system that allows for dynamical representations, and long-term memory. These requirements are not novel in cognitive modeling, and in fact constitute some of the canons of the connectionist framework, in which, an agents internal state is often described as a trajectory of activity patterns influenced by attractors in recurrent networks. Long-term memory is supported by adaptations to the weights or connectivity of the network, which gradually evolve as the agents internal state interacts with various learning algorithms. Such connectionist modelling has been applied successfully in recent developmental robotics projects with a social learning component, e.g. [3] endowed the iCub with a connectionist architecture to mirror and explain body-centered learning effects in infant/caregiver interactions.

References:

[1] Hannagan, T. and Wu, R. (2011) “Cued multimodal learning in infancy: A neuro-computational model,” In C. Hoelscher, T. F. Shipley and L. Carlson (Eds.), Proceedings of the 33rd Annual Conference of the Cognitive Science Society. Boston, MA: Cognitive Science Society.
[2] Wu, R. and Kirkham, N. Z. (2011) “No two cues are alike: Depth of learning during infancy is dependent on what orients attention,” Journal of Experimental Child Psychology, vol. 107, pp. 118–136.
[3] Morse, A. F., De Greeff, J., Belpaeme, T. and Cangelosi, A. (2010) “Epigenetic robotics architecture (era),” IEEE Transactions on Autonomous Mental Development, vol. 2, no. 4, pp. 325–339.

A EEG Based Human-Robot Interaction: Implications for Developmental Robotics

Stevo Bozinovski
Department of Mathematics and Computer Science, South Carolina State University, USA

As pointed out by Rohlfing and Wrede ([6]) the current methods of developmental robotics are mostly based on modeling children’s development and designing various types of learning systems. They point attention to Human Robot Interaction (HRI). They also point out that it is not only the input but rather the interplay between the input and the feedback, as well as modality of the feedback (e.g. eye gaze), that enables the learner to take advantage in the learning process [5].

In this discussion we propose a possible direction in developmental robotics, the *EEG based HRI*. We describe a paradigm that involves learning, and on which we have been working for some time.

The EEG based HRI started in 1988 when the first control of a robot using EEG signal was reported [1], followed by the robot control using EOG signals [2]. Shortly after that taxonomy of brain potentials was proposed which explicitly introduces the class of anticipatory brain potentials [3], an experimental paradigm was carried out exploring learning and development in the brain observed as a cognitive wave named Electroexpectogram (EXG). Recently the paradigm was used in controlling two robotic arms toward solving a three-disk Tower of Hanoi task [4].

The EXG is obtained in a paradigm named *CNV flip-flop paradigm*. It is a feedback version of the classical Contingent Negative Variation (CNV) paradigm. The paradigm involves a naïve subject (learner) who was only told to press a button whenever the distinct S2 stimulus appears. The learner quickly realizes that there is an association between S1 and S2 and develops expectancy *expect*(S2/S1) which elicits a CNV potential. The CNV potential, which appears between the stimuli S1 and S2, has a distinctive ramp-shape. At that point, the *interaction partner* (e.g. a robot), disables the S2 stimulus and performs some other action, for example moving its robotic arm. As experiment progresses, the human subject realizes that there is no more association between S1 and S2, which lowers the *expect(S2/S1)* down to the point where the observed event related potential has no CNV shape. Now the interaction partner turns on the stimulus S2 again (and optionally moves the robotic arm), and so on. The subject keeps adapting to the changing environment, which produces oscillation of *expect(S2/S1)*. The oscillation of a CNV parameter (e.g. regression slope) is presented as the Electroexpectogram (EXG) curve.

When the experiment ends, the EXG is a record of mental development of a human brain in respect to the CNV flip-flop environment. And it was stored in the robot “brain” interactively, during interplay with the human brain. The EXG provides a learned experience which can drive a sequence of future actions. So, the CNV flip-flop paradigm is a multimodal HRI paradigm in which a robot mental development is obtained in an interactive process with a human EEG.

References:

[1] Bozinovski S., Sestakov M., Bozinovska L. (1988) Using EEG Alpha Rhythm to Control a Mobile Robot", In G. Harris, C. Walker (eds.) *Proc 10th Annual Conf. of the IEEE Engineering in Medicine and Biology Society,* Vol 3, pp. 1515-1516
[2] Bozinovski S. (1990) Mobile Robot Trajectory Control: From Fixed Rails to Direct Bioelectric Control, in O. Kaynak (Ed.) *Proc IEEE International Workshop on Intelligent Motion Control*, vol 2, pp. 463-467
[3] Bozinovska L., Bozinovski S., Stojanov G., (1992) Electroexpectogram: Experimental Design and Algorithms, Proc. IEEE International Biomedical Engineering Days, pp. 58-60
[4] Bozinovski A., Tonkovic S., Isgum V., Bozinovska L (2011) Robot Control Using Anticipatory Brain Potentials. *Automatika 52* (1). 20-30
[5] Pitsch, K., Vollmer, A.-L., Fritsch, J., Wrede, B., Rohlfing, K. J. & Sagerer, G. (2009) On the loop of action modification and the recipient's gaze in adult-child interaction. In: *Proceedings of the International Conference on Speech and Gesture in Interaction, September (GESPIN).*
[6] Rohlfing K. and Wrede B.(2012) What novel scientific and technological questions developmental robotics bring to HRI? – Are we ready for a loop?, *AMD Newsletter 8* (2). 11-13

The Developmental Drive: From Designing Control Towards Enabling Interaction

Helge Ritter
Faculty of Technology & Excellence Cluster "Cognitive Interaction Technology" (CITEC)
Bielefeld University, Germany

A major temptation in robotics – driven by a desire to rapidly create useful applications – has been to "aim high" and focus on tasks that are very different from what we are concerned with during our early months or even years of life (since during that phase our primary goal is not exactly to be useful).

The connection between developmental psychology and robotics has made us appreciate that there is a rich fabric of sensorimotor and social skills acquired by a child in the first months of its life. This has begun to transform the field of robotics, triggering projects emphasising so far largely neglected aspects in robotics, such as the role of changing body morphologies [1,2], the availability of a sensitive skin surrounding the whole body [3] or developmental strategies for integrating multiple capabilities and skills [4].

The area where the benefits of a deeper understanding of developmental aspects for robotics already begins to become particularly clear is learning, giving us valuable clues how to organize suitable "scaffolds" or detect "contingencies" [5] to guide low-level sensory learning towards the high-level learning that humans or higher animals exhibit even in very complex environments. Besides the important social level, the body morphology as a major developmental scaffold is about to trigger new research, such as focusing on the role of different body regions in the organization of mental growth and the development of cognition [3]. A major example focused in our own work is "manual intelligence" [6], the cluster of abilities developing around our use of hands.

So we are definitely ready for the loop. What can we expect from it? My hunch is that developmental cognition research will drag us away from a style of algorithm development using handcrafted representations based primarily on nameable, symbolic entities and makes us ready for the challenge of creating interfaces and adaptation rules coping with the interactive shaping of dynamic processes with strong holistic and subsymbolic traits. This may give a new gist to our science, developing a body of knowledge how to reproducibly set up the right conditions for guiding processes of dynamic growth towards cognitive and skill development in artificial systems, without having to mess with the details themselves. And it may shift our taste and attitude of understanding from a design of static control frameworks towards dynamic architectures for realizing adaptive, rich and ressource-sensitive cognitive interaction.

References:

[1] Bongard, J. (2011) Morphological change in machines accelerates the evolution of robust behavior. PNAS 108(4):1234-1239

[2] Kaplan, F., Oudeyer, P.Y (2009) Stable kernels and fluid body envelopes SICE Journal of Control, Measurement and System Integration 48(1)

[3] Asada et al. (2009) Towards Computational Developmental Models based on Synthetic Approaches IEEE 8th Int. Conf. on Development and Learning pp. 1-8

[4] Cangelosi, A. et al. (2010) Integration of Action and Language Knowledge: A Roadmap for Developmental Robotics IEEE Autonomous Mental Development 2(3) 167-195

[5] Lohan, K. et al (2011) Contingency allows the robot to spot the tutor to learn from interaction. IEEE Int. Conf. on Development and Learning, pp. 1-8

[6] Ritter, H., Haschke, R., Röthling, F. and Steil, J.J. (2011) Manual Intelligence as a Rosetta Stone for Robot Cognition ROBOTICS RESEARCH / Springer Tracts in Advanced Robotics 66:135-146

Reply and Summary:
A loop takes time

Katharina J. Rohlfing & Britta Wrede

Center of Excellence Cognitive Interaction Technology (CITEC) Bielefeld University, Germany

The discussion shows that two perspectives on social interaction need to be differentiated. From one perspective, and as pointed out by Lakshmi Gogate as well as Anthony Morse, Thomas Hannagan and Rachel Wu, social interaction constitutes a source of cues that can effectively guide attention. Its effects can be viewed as comparable to other sources that guide attention as well [10]. From the other perspective, however, in order to design interactions that can take advantage of social cues, we need to understand how infants accumulate experiences from social interaction, i.e. which experiences they need, to later understand and appropriately take advantage of social cues. For this, it is necessary to consider the "history of social interaction" ([7], p. 192) — this is why we think that a loop should not only involve the learner's side but also the history of the interaction between the social partners. We believe that – incorporating insights from developmental studies – Developmental Robotics bring three important strings of research to the HRI:

Firstly, HRI research can assume that the attention of a learner can be guided ([19]) to the relevant aspects of an event, which can then be picked up. From current studies, we gain lots of insights into the ways in which contextual information is important in this attention guiding process; for example what is new in a particular situation ([14]) as well as on-line lexical competition among activated words and their meanings [22]. Yet we know little about how these cues are embedded in social interaction. Some insights are provided by studies which – as Lakshmi Gogate is pointing out – are combining observational and experimental techniques. From these studies, we know for example that to link a label onto a referent, young learners need to be provided with synchronous movement ([8]) and that parents provide such scaffolding for younger but not for older children [9]. However, these approaches can only tell us about the fast attribution of new word and object as well as about the strategies that are necessary for a child to form this initial link. From recent research in word learning ([13]), however, we know that this link is fragile and needs to be strengthened by repeated exposure. Thus, the question is: What happens when a learning situation is repeated in order to facilitate long-term learning? Is it simply sufficient to initialize the same strategy (e.g. to provide the child with the label again and again while moving the object) or is it rather necessary to change the strategy and to fine-tune to the learner's understanding? Current research suggests that while in the first iteration, directing children's attention to the referent is crucial, in the second iteration, it might be advantageous to direct attention away from competitors [3].

Secondly, we need to understand how infants' attention becomes educated towards social and contextual cues. According to the Emergentist Coalition Model ([10], [12]), it seems that from a particular age on, children are more sensitive towards social cues than before [17]. Even though the authors explain this developmental change by pointing to the possibility that "infants come to recognize people as intentional beings who have goals" ([10], p. 32), together with Anthony Morse, Thomas Hannagan and Rachel Wu, we think that, again, the history of interaction is responsible for the effective presence of social cues in videos. Accordingly, the observation that learning of audio-visual events improves when presented with social cues might highly depend on "previous appropriate feedback experience". Our own research lets us suggest that sensitivity towards social cues does not emerge over night. Even though the understanding of pointing can be attested to children around their first birthday, we have shown that in 4.5 month-olds, a rudimentary understanding of pointing can be observed under specific conditions [18]. Thus, children are capable of understanding a referential gesture early in their development, when it is supported by a dynamic movement.

These rudimentary capabilities are not innate but have to be acquired (e.g. [4]). Thus, we agree with Stevo Bozinovski claiming that it is important to create expectations about situational outcomes, and EEG research with young infants shows that they create expectancy in the form of behavioral disposition before they produce this behavior [11]. Yet, we still know little about how such expectations are built. In the research by Flom and colleagues, for example, it is suggested that "infants might learn to follow gaze or points to visually accessible locations because most triadic interactions in the first year involve nearby, visible targets" (p. 192). For HRI, the unfolding expectation implies that knowledge is not binary coded but might have different forms varying from weakly activated to decontextualized. Thus, even though Adriana Tapus is suggesting to make artificial systems innately sensitive to social cues, we think that in the ecology of learning, different sensitivity to these cues might evoke different forms of learning and knowledge — maybe necessary on the way to proficiency. Along these lines, research that Yukie Nagai is conducting promise to reveal more about how different levels of body coordination of the learner shape the movement complexity of the tutor. Yet, we need to understand not only what the channels are, through which such information is flowing depending on the interlocutor's abilities ([5]) but also how these channels dynamically switch during an ongoing interaction to make it more efficient [21]. This is a feat accomplished by Andrea Thomaz and colleagues in robot teaching ([1-2]). In this work, tutors switched between teaching holistic trajectories and (static) keyframes. This behavior makes it possible for the tutor to focus on relevant aspects of learning in reaction to the learner's understanding. The authors are thus paving the way towards a much more flexible way of semantic-oriented robot teaching which is also postulated by Ritter who sees the need to combine dynamic processes with holistic traits, in order to move away from handcrafted representations as they are currently used in robotics.

Thirdly, we can assume that social cues foster long-term learning. Taking, again, examples from word learning studies, Horst & Samuelson (2008) have shown that words are better retained when provided within an ostensive interaction. Similarly, the results in McGregor and colleagues (2009) suggest that words are learned better when a gesture is shown to the learner as opposed to a picture. Thus, referential gestures seem not only to have an impact on the online task solution but also on the retention and learning from the situation. However, we think that ostensive cues can only be created when taking into account the history of interaction. Calling a child's name is effective only when this prosodic pattern deviates from what the child heard right before. Thus, again, it is not a particular behavior (i.e. calling somebody's name) that is ostensive per se, but it is this signal put into a particular communicative context that creates an ostensive effect making particular information relevant [20]. We think that this idea can be pushed even further: A tutor can make a whole interaction relevant to the child by behaving contingently with the learner, as also suggested in the response by Yukie Nagai. We think that in this way, pragmatic frames emerge that "tell" the child in which spot important information can be found. Such a pragmatic frame can be observed when, for example, a new referent is labeled via pointing; the whole act of labeling consists then of eye contact, ostensive speech and the pointing gesture directed at the learner — all together forming a pattern, within which a child can rapidly pick up a word for a referent. How such patterns are established, is still not well explained. Yet, design for human-robot interaction can benefit from the identification of such patterns ([6], [15]), enabling systems to perceive interaction more globally (as patterns) rather than being overloaded with parsing every single unit. Also, robot teaching scenarios could benefit from the use of natural, possibly generalizing, cues instead of making use of programming commands to distinguish between different teaching styles (as e.g. in [2]).

Overall, we not only share the view of Ritter and Thomaz that developmental robotics is a promising approach to build more flexible resource-sensitive robots, but that it is a unique research area which is very creative in developing new methods for analyzing HOW capabilities can incrementally emerge — in robots as well as in infants, to facilitate intuitive interaction.

References:

[1] Akgun, B., Cakmak, M., Woo Yoo, J., Tomaz, A. L. (2012) Trajectories and keyframes for kinesthetic teaching: a human-robot interaction perspective. In: Proceedings of the Human-Robot Interaction (ACM-IEEE HRI): 391–398.
[2] Akgun, B., Jiang, K., Cakmak, M. & Thomaz, A. (2011) Learning tasks and skills together from a human teacher. In: Proceedings of the Twenty-Fifth AAAI Conference on Artificial Intelligence: 1868–1869.
[3] Axelsson, E. L., Churchley, K. & Horst, J. S. (2012) The right thing at the right time: why ostensive naming facilitates word learning. In: Frontiers in Psychology 3: Article 88.
[4] Daum, M. M., & Gredebäck, G. (2011) The development of grasping comprehension in infancy: Covert shifts of attention caused by referential actions. In: Experimental Brain Research, 208: 297–307.
[5] Fischer, K., Foth, K., Rohlfing, K. J., & Wrede, B. (2011) Mindul tutors: Linguistic choice and action demonstration in speech to infants and robots. In: Interaction Studies 12: 134–161.
[6] Fischer, K. (2010) Why it is interesting to investigate how people talk to computers and robots: Introduction to the special issue. In: Journal of Pragmatics 42: 2349–´2354.
[7] Flom, R., Deák, G. O., Phill, C. G. & Pick, A. D. (2004) Nine-month-olds' shared visual attention as a function of gesture and object location. In: Infant Behavior & Development 27: 181–194.
[8] Gogate, L. J. & Bahrick, L. E. (2001) Intersensory redundancy and 7-month-old infants' memory for arbitrary syllable-object relations. In: Infancy 2: 219–231.
[9] Gogate, L., Bahrick, L. & Watson, J. (2000) A study of multimodal motherese: The role of temporal synchrony between verbal labels and gestures. In: Child Development 71: 878–894.
[10] Golinkoff, R. M., & Hirsh-Pasek, K. (2006) Baby wordsmith: From associationist to social sophisticate. In: Current Directions in Psychological Science 15: 30–33.
[11] Gredebäck, G., Melinder, A., & Daum, M. (2010) The development and neural basis of pointing comprehension. In: Social Neuroscience 5: 441–450.
[12] Hollich, G., Hirsh-Pasek, K., & Golinkoff, R. (2000) Breaking the language barrier: An emergentist coalition model of word learning. Monographs of the Society for Research in Child Development 65.
[13] Horst J. & Samuelson L. (2008), Fast Mapping But Poor Retention in 24-Month-Old Infants. In: Infancy, Volume 13: 128–157
[14] Kucker, S. C. & Samuelson, L. K. (2011) The first slow step: Differential effects of object and word-form familarization on retention of fast-mapped words. In: Infancy: 1–29
[15] Lohse, M. (2011) The role of expectations and situations in human-robot interaction. In: New Frontiers in Human-Robot Interaction: John Benjamins Publishing Company: 35–56.
[16] McGregor, K. K., Rohlfing, K. J., Bean, A. & Marschner, E. (2009) Gesture as a support for word learning: The case of under. In: Journal of Child Language 36: 807–828.
[17] Pruden, S. M., Hirsh-Pasek, K., Michnick Golinkoff, R. & Hennon, E. A. (2006) The birth of words: Ten-month-olds learn words through perceptual salience. In: Child Development 77: 266–280.
[18] Rohlfing, K. J., Longo, M. R. & Bertenthal, B. I. (2012) Dynamic pointing triggers shift of visual attention in young infants. In: Developmental Science 15: 426–435.
[19] Smith, L. B., Colunga, E. & Yoshida, H. (2010) Knowledge as process: contextually cued attention and early word learning. In: Cognitive Science 34: 1287–1314.
[20] Sperber, D. & Wilson, D. (1995) Relevance: Communication and Cognition. Oxford: Blackwell.
[21] Wrede, B., Kopp, S., Rohlfing, K. J., Lohse, M. & Muhl, C. (2010) Appropriate feedback in asymmetric interactions. In: Journal of Pragmatics 42: 2369–2384.
[22] Yoshida, H. & Hanania, R. (2011) If it's red, it's not Vap: how competition among words may benefit early word learning. In: First Language: 1–17.

Dialog Initiation

Children's Natural Learning: Why development really does matter!

Denis Mareschal
Centre for Brain and Cognitive Development
School of Psychology
Birkbeck College
University of London, London, UK.

There have been many significant advances in the field of developmental cognitive neurosciences over the last few years. These have important implications for the way learning in developing systems is to be considered. To date these factors are rarely directly considered in the existing (virtual or robotic) models of development. I will list three of these and suggest that they pose significant challenges for current computational and robotic models of learning in developing systems.

Epigenetics: When genes are turned on and off

One simple view of development is that it is the outcome of genetic "innate" constraints and environmental "experiential" constraints. Indeed, the field of *Behavioural Genetics* is largely concerned with trying to partition any observed variation in behaviour into a "genetic component" and an "environmental component" [8]. A more complex view is that these processes interact. So, for example, a constant genetic predisposition may have a greater or lesser impact on development depending on the agent's environment [11]. *In fact, the situation is far more complex.* It is now clear that gene expression itself can be self-modulated depending on the environment ([4-5], [7]). For example, environmental changes such as an absence of food can lead to brain chemical imbalances in worker bees that alter the expression of genes, and consequently the physiological and functional roles of these bees in a hive. In other words… the effective genetic constraints are not constant and depend on environmental pressures. The extent to which genetic material is expressed depends on the environmental needs of the agent. This is not only true in bees, but may play an important role in the expression of complex cognitive behaviors [3].

Morphogenesis: body growth does matter in early learning

The brain is particularly plastic during the early years [10]. This is not just so that children can acquire new knowledge easily, but is also true because there is a need to constantly re-calibrate sensory information in a sensory-motor system that is dramatically changing in size and sensory efficacy. Indeed, while adults are able to combine sensory cues optimally to improve sensory estimates, children do not appear to do so until 8-12 years of age (e.g., [6]). This is because child's changing physical dimensions (e.g., separation eye) continually distort the possible interpretation of sensory input. Body size also acts as an effective filter on the complexity of the environment children learn from. For example, arm length helps support what they are attending to because objects closer up will block larger portion of visual field [9].

Affect and Trust: Not all teachers are equal

It is now well established that social interactions form an important part of how children learn [12]. In particular, children can only acquire some knowledge (such as the existence of germs) through the testimony of others and not through direct experience. However, children do not learn equally from all social interactions. In fact, from the earliest ages children identify those adults or peers in whose testimony they can trust [2]. This often leads to increased attachment and affect for that person. At the neural level, increased positive affect leads to the releases of dopamine throughout key parts of the brain that has the consequence of increasing plasticity in those parts of the brain [1]. Thus, positive affect plays a role in modulating learning both at the neural level and at the (macro) social level.

These three factors (among others) result in an effective learning environment that is highly adaptive to the current needs of the learner. Importantly, it is a very different environment for a developing agent than for the fully developed, adult agent. Traditionally, computational modelers have tended to characterize learning systems in terms of the mechanisms and processes present in the adult. My claim is that they also need to recognize the unique character of learning in a truly developmental system.

References:

[1] Dayan, P & Huys, Q. J. (2009) Serotonin in affective control. *Annual Review of Neuroscience, 32*, 95-126

[2] Harris, P. L., Pasquin, E.S., Duke, S., Asscher, J. & Pons, F. (2005) Germs and angels: The role of testimony in children's ontology. *Developmental Science, 9*, 76-96.

[3] Kan, K. J., Ploeger, A., Raijmakers, M. E. J., Dolan, C. V., & van der Maas, H. L. J. (2010). Nonlinear Epigenetic Variance: Review and Simulations. *Developmental Science, 13(1)*, 11-27.

[4] Mareschal, D. Johnson, M. H., Sirois, S., Spratling, M., Thomas, M. & Westermann, G. (2007) *Neuroconstuctivism Vol. 1: How the brain constructs cognition.* Oxford UK: OUP.

[5] Meaney, M. J. (2010) Epigenetics and the Biological Definition of Gene × Environment Interactions. *Child Development, 81*, 41-79.

[6] Nardini, M., Bedford, R. & Mareschal, D (2010). Fusion of visual cues is not mandatory in children. *Proceedings of the National Academy of Science U. S. A. 107*, 17041-17046

[7] Ronald, A. (2011). Is the Child father of the man? Evaluating the stability of genetic influences across development. *Developmental Science, 14*, 1471-8.

[8] Rutter, M. (2008) *Genes and Behavior*. Oxford, UK: Blackwell Publihsing

[9] Smith, L. B., Yu, C., & Pereira, A. F. (2011) Not your mother's view: the dynamics of toddler visual experience. *Developmental Science, 14,* 9-17.

[10] Thomas, M.S.C. and Johnson, M. H. (2006). The computational modelling of sensitive periods. *Developmental Psychobiology*, *48(4)*, 337-344.

[11] Turkheime, E., Haley, A., Waldron, M, D'Onofrio, B. & Gottesman, I. I. (2003) Socioeconomic status modifies heritability of IQ in Young Children. *Psychological Science, 14*, 623-628.

[12] Vygotsky (1976) *Mind in Society*. Harvard University Press.

Call for paper and tutorials
IEEE International Conference on Development and Learning and Epigenetic Robotics (ICDL-EpiRob), 2012

Paper submission deadline: June 15, 2012
Tutorial submission deadline: June 15, 2012
Notification of acceptance: September 15, 2012
Conference: November 7-9, 2012

Location: San Diego, California, USA
Web: http://www.icdl-epirob.org/
General chairs: Javier Movellan (UCSD, US), Matthew Schlesinger (SIU Carbondale, US),
Jochen Triesch (FIAS, Germany)
Program Chairs: Yukie Nagai (Osaka Univ. Japan), Ian Fasel (Univ. Arizona, US), Clay Morrison (Univ. Arizona, US)

The ICDL and the Epigenetic Robotics conferences are the premier venues for interdisciplinary research that blends the boundaries between robotics, artificial intelligence, machine learning, developmental psychology, and neuroscience. The scope of development and learning covered by this conference includes perceptual, cognitive, motor, behavioral, emotional and other related capabilities that are exhibited by humans, higher animals, artificial systems and robots.

Topics of interest include – but are not limited to:
• The science and engineering of learning.
• Machine learning and development.
• Sensory motor control.
• Computational approaches to the study of development and learning.
• Neural basis of development and learning.
• Language Acquisition.
• Neurogenesis.
• Datasets for the Analysis of Learning and Development.
• Field Studies.
• Biomimetic robots.
• Affect, Development and Learning.
• Intrinsic motivation, exploration and play
• Social development in humans and robots.
• Applications to education and clinical interventions.

In addition to article submissions, experts in different areas are invited to organize a 3-hour tutorial, which will be held on the first day of the conference. Participants in tutorials are asked to register for the main conference as well. Tutorials are meant to provide insights into specific topics as well as overviews that will inform the interdisciplinary audience about the state-of-the-art in child development, neuroscience, robotics, or any of the other disciplines represented at the conference.

Call for Participation
ICIS 2012 Developmental Robotics Workshop

Date: June 14, 2012, 1:00-6:00 PM
Location: Minneapolis, Minnesota, USA

Web: http://www.frontiersin.org/events/ICIS_2012_Preconference_Workshop_on_Developmental_Robotics/1663

The purpose of this preconference workshop is to introduce behavioral researchers to the emerging field of developmental robotics. In particular, 8 invited speakers will provide an accessible overview of the robotic platforms, computational algorithms, and experimental methods used in developmental robotics. The workshop is part of a long-term effort to establish and support collaborative research that bridges the study of learning and development across both natural and artificial systems. As a result, the International Conference on Infant Studies (i.e., ICIS, the largest organization for the study of human infants) was selected as the site for the workshop.

The workshop is organized on behalf of the IEEE-CIS AMD Technical Committee by Matthew Schlesinger and Jochen Triesch. The panel of guest speakers includes: Minoru Asada (Osaka University), Giorgio Metta (ITT, Italy), Anthony Morse (University of Plymouth), Javier Movellan (UC San Diego), Katharina Rohlfing (University of Bielefeld), Brian Scassellati (Yale University), Matthew Schlesinger (SIU Carbondale).

BMI Summer School
International Conference on Brain-Mind (ICBM)

BMI Summer School: June 25 - August 3, 2012
International Conference on Brain-Mind: July 14 - July 15, 2012
Location: Michigan State University, East Lansing, Michigan, USA

Web: http://www.brain-mind-institute.org/

Collectively, the human race seems ready to unveil one of its last mysteries - how its brain-mind works at computational depth. The research community needs a large number of leaders who have sufficient knowledge in at least six disciplines conjunctively - Biology, Neuroscience, Psychology, Computer Science, Electrical Engineering, and Mathematics (6 disciplines). The Brain-Mind Institute (BMI) provides an integrated 6-discipline academic and research infrastructure for future leaders of brain-mind research. The BMI is a new kind of institute, not limited by boundaries of disciplines, organizations, and geographic locations.

The subjects of interest for ICBM include, but not limited to:
Genes, Cells, Circuits, Diseases.
Streams: pathways, attention, vision, audition, touch, taste.
Brain ways: neural networks, brain-mind architecture, inter-modal, neural modulation.
Experiences/learning: training, learning, development, interaction, intelligence metrics.
Behaviors: actions, concept learning, abstraction, languages, decision, reasoning.

Societies/multi-agent: joint attention, swarm intelligence, group intelligence, laws.
Applications: image analysis, computer vision, speech recognition, pattern recog.

The BMI summer courses 2012:
BMI 811 Biology for Brain-Mind Research, June 25 - July 13, 2012.
BMI 821 Neuroscience for Brain-Mind Research, July 16 - Aug.3, 2012.
BMI 871 Computational Brain-Mind (distance learning): Aug. 6 - Aug. 24, 2012

Keynote talks:
James L. McClelland, Stanford University
Stephen Grossberg, Boston University
Jim L. Olds, George Mason University

Call for paper: Third International Workshop on Human-Behavior Understanding (HBU 2012) to be held in conjunction with IROS 2012

Paper submission deadline: July 1st, 2012
Notification of acceptance: July 25th, 2012
Workshop: October 7-12, 2012
Location: Vilamoura, Algarve, Portugal
Organizing committee: Albert Ali Salah (Bogazici Univ., Turkey), Javier Ruiz-del-Solar (Univ. Chile, Chile), Cetin Mericli (CMU, US), Pierre-Yves Oudeyer (Inria, France).

Web: http://www.cmpe.boun.edu.tr/hbu/2012/

The Third Workshop on Human Behavior Understanding, organized as a satellite to IROS'2012, will gather researchers dealing with the problem of computational modeling and understanding of human behavior under its multiple facets (expression of emotions, display of relational attitudes, performance of individual or joint actions, imitation, etc.), with particular attention to implications in robotics, including additional resource and robustness constraints of robotic platforms, social aspects of human-robot interaction, and developmental approaches to robotics.

The HBU Workshops, previously organized as satellite to ICPR and AMI Conferences, have a unique aspect of fostering cross-pollination of different disciplines, bringing together researchers of robotics, HCI, artificial intelligence, pattern recognition, interaction design, ambient intelligence, psychology. The diversity of human behavior, the richness of multi-modal data that arises from its analysis, and the multitude of applications that demand rapid progress in this area ensure that the HBU Workshops provide a timely and relevant discussion and dissemination platform.

IEEE TRANSACTIONS ON AUTONOMOUS MENTAL DEVELOPMENT

Volume 3, Issue 4, December 2011
Link: http://ieeexplore.ieee.org/xpl/tocresult.jsp?isnumber=6097099&punumber=4563672

Guest Editorial: Special Issue on Computational Modeling of Neural and Brain Development
Jin, Y. N.; Meng, Y. M.; Weng, J. W.; Kasabov, N. K. (PDF)

From Infant Brains to Robots: A Report From the IEEE International Conference on Development and Learning (ICDL)-International Conference on Epigenetic Robotics (EpiRob) 2011 Conference
Cangelosi, A; Triesch, J. (PDF)

A Model of Neuronal Intrinsic Plasticity

Chunguang Li. Page(s): 277 - 284 (PDF)
Abstract: Recent experimental results have accumulated evidence that the neurons can change their response characteristics to adapt to the variations of the synaptic inputs, which is the so-called neuronal intrinsic plasticity mechanism. In this paper, we present a new model on neuronal intrinsic plasticity. We first show that the probability distribution of the neuronal firing rates is more suitable to be represented as a Weibull distribution than an exponential distribution. Then, we derive the intrinsic plasticity model based on information theory. This study provides a more realistic model for further research on the effects of intrinsic plasticity on various brain functions and dynamics.

Firing Rate Homeostasis for Dynamic Neural Field Formation

Glaser, C.; Joublin, F. Page(s): 285 - 299 (PDF)
Abstract: Dynamic neural fields are recurrent neural networks which aim at modeling cortical activity evolution both in space and time. A self-organized formation of these fields has been rarely explored previously. The main reason for this is that learning-induced changes in effective connectivity constitute a severe problem with respect to network stability. In this paper, we present a novel network model which is able to self-organize even in face of experience-driven changes in the synaptic strengths of all connections. Key to the model is the incorporation of homeostatic mechanisms which explicitly address network stability. These mechanisms regulate activity of individual neurons in a similar manner as cortical activity is controlled. Namely, our model implements the homeostatic principles of synaptic scaling and intrinsic plasticity. By using fully plastic within-field connections our model further decouples learning from topological constraints. For this reason, we propose to incorporate an additional process which facilitates the development of topology preserving mappings. This process minimizes the wiring length between neurons. We thoroughly evaluated the model using artificial data as well as continuous speech. Our results demonstrate that the network is able to self-organize, maintains stable activity levels, and remains adaptive to variations in input strength and input distribution.

Probabilistic Computational Neurogenetic Modeling: From Cognitive Systems to Alzheimer's Disease

Kasabov, K.; Schliebs, R.; Kojima, H. Page(s): 300 - 311 (PDF)
Abstract: The paper proposes a novel research framework for building probabilistic computational neurogenetic models (pCNGM). The pCNGM is a multilevel modeling framework inspired by the multilevel information processes in the brain. The framework comprises a set of several dynamic models, namely low (molecular) level models, a more abstract dynamic model of a protein regulatory network (PRN) and a probabilistic spiking neural network model (pSNN), all linked together. Genes/proteins from the PRN control parameters of the pSNN and the spiking activity of the pSNN provides feedback to the PRN model. The overall spatio-temporal pattern of spiking activity of the pSNN is interpreted as the highest level state of the pCNGM. The paper demonstrates that this framework can be used for modeling both artificial cognitive systems and brain processes. In the former application, the pCNGM utilises parameters that correspond to sensory elements and neuromodulators. In the latter application a

pCNGM uses data obtained from relevant genes/proteins to model their dynamic interaction that matches data related to brain development, higher-level brain function or disorder in different scenarios. An exemplar case study on Alzheimer's Disease is presented. Future applications of pCNGM are discussed.

A Multiple Context Brain for Experiments With Robot Consciousness

Andreae, J.H. Page(s): 312 - 323 (PDF)

Abstract: The PURR-PUSS system (PP) is a versatile model of a human-like brain, designed to be implemented in parallel hardware and embodied in the head of a robot moving in the real world. The aim of the research with PP is to try out mechanisms for learning, intelligence and consciousness. Limitations of resources have dictated that the experiments with PP are made on a personal computer by simulating the brain and robot body in a microworld. The unique features of PP are multiple context and novelty-seeking. In this paper, a squash-pop microworld is described first, so that concrete examples can be given for a brief review of the PP system, followed by two new features called trail memory, to realize Baars' global workspace, and belief memory, to realize Rosenthal's higher order thoughts and Johnson-Laird's conscious reasoning. The extended system, PP*, is designed to give consciousness to the subconscious PP, but higher order thoughts and conscious reasoning prove to be elusive. A definition of a conscious robot provides a measure of progress.

Volume 4, Issue 1, March 2012

Link: http://ieeexplore.ieee.org/xpl/tocresult.jsp?isnumber=6167349&punumber=4563672

Episodic-Like Memory for Cognitive Robots

Stachowicz, D.; Kruijff, G.M. Page(s): 1 - 16 (PDF)

Abstract: The article presents an approach to providing a cognitive robot with a long-term memory of experiences-a memory, inspired by the concept of episodic memory (in humans) or episodic-like memory (in animals), respectively. The memory provides means to store experiences, integrate them into more abstract constructs, and recall such content. The paper presents an analysis of key characteristics of natural episodic memory systems. Based on this analysis, conceptual and technical requirements for an episodic-like memory for cognitive robots are specified. The paper provides a formal design that meets these requirements, and discusses its full implementation in a cognitive architecture for mobile robots. It reports results of simulation experiments which show that the approach can run efficiently in robot applications involving several hours of experience.

A Model to Explain the Emergence of Imitation Development Based on Predictability Preference

Minato, T.; Thomas, D.; Yoshikawa, Y.; Ishiguro, H. Page(s): 17 - 28 (PDF)

Abstract: Imitation is a very complicated function which requires a body mapping (a mapping from observed body motions to motor commands) that can discriminate between self motions and those of others. The developmental mechanism of this sophisticated capability, and the order in which the required abilities arise, is poorly understood. In this paper, we present a mechanism for the development of imitation through a simulation of infant-caregiver interaction. A model was created to acquire a body mapping, which is necessary for successful mutual imitation in infant-caregiver interaction, while discriminating self-motion from the motion of the other. The ability to predict motions and the time delay between performing a motion and observing any correlated motion provides clues to assist the development of the body mapping. The simulation results show that the development of imitation capabilities depends on a predictability preference (a function of how an agent feels regarding its options of "what to imitate," given its ability to predict motions). In addition, the simulated infants in our system are able to develop the components of a healthy body mapping in order, that is, relating self motion first, followed by an understanding of others' motions. This order of development emerges spontaneously without the need for any explicit mechanism or any partitioning of the interaction. These results suggest that this predictability preference is an important factor in infant development.

Symbolic Models and Emergent Models: A Review

Juyang Weng. Page(s): 29 - 53 (PDF)

Abstract: There exists a large conceptual gap between symbolic models and emergent models for the mind. Many emergent models work on low-level sensory data, while many symbolic models deal with high-level abstract (i.e., action) symbols. There has been relatively little study on intermediate representations, mainly because of a lack of knowledge about how representations

fully autonomously emerge inside the closed brain skull, using information from the exposed two ends (the sensory end and the motor end). As reviewed here, this situation is changing. A fundamental challenge for emergent models is abstraction, which symbolic models enjoy through human handcrafting. The term abstract refers to properties disassociated with any particular form. Emergent abstraction seems possible, although the brain appears to never receive a computer symbol (e.g., ASCII code) or produce such a symbol. This paper reviews major agent models with an emphasis on representation. It suggests two different ways to relate symbolic representations with emergent representations: One is based on their categorical definitions. The other considers that a symbolic representation corresponds to a brain's outside behaviors observed and handcrafted by other outside human observers; but an emergent representation is inside the brain.

A Behavior-Grounded Approach to Forming Object Categories: Separating Containers From Noncontainers

Griffith, S.; Sinapov, J.; Sukhoy, V.; Stoytchev, A. Page(s): 54 - 69 (PDF)
Abstract: This paper introduces a framework that allows a robot to form a single behavior-grounded object categorization after it uses multiple exploratory behaviors to interact with objects and multiple sensory modalities to detect the outcomes that each behavior produces. Our robot observed acoustic and visual outcomes from six different exploratory behaviors performed on 20 objects (containers and noncontainers). Its task was to learn 12 different object categorizations (one for each behavior-modality combination), and then to unify these categorizations into a single one. In the end, the object categorization acquired by the robot matched closely the object labels provided by a human. In addition, the robot acquired a visual model of containers and noncontainers based on its unified categorization, which it used to label correctly 29 out of 30 novel objects.

Autonomous Learning of High-Level States and Actions in Continuous Environments

Mugan, J.; Kuipers, B. Page(s): 70 - 86 (PDF)
Abstract: How can an agent bootstrap up from a low-level representation to autonomously learn high-level states and actions using only domain-general knowledge? In this paper, we assume that the learning agent has a set of continuous variables describing the environment. There exist methods for learning models of the environment, and there also exist methods for planning. However, for autonomous learning, these methods have been used almost exclusively in discrete environments. We propose attacking the problem of learning high-level states and actions in continuous environments by using a qualitative representation to bridge the gap between continuous and discrete variable representations. In this approach, the agent begins with a broad discretization and initially can only tell if the value of each variable is increasing, decreasing, or remaining steady. The agent then simultaneously learns a qualitative representation (discretization) and a set of predictive models of the environment. These models are converted into plans to perform actions. The agent then uses those learned actions to explore the environment. The method is evaluated using a simulated robot with realistic physics. The robot is sitting at a table that contains a block and other distractor objects that are out of reach. The agent autonomously explores the environment without being given a task. After learning, the agent is given various tasks to determine if it learned the necessary states and actions to complete them. The results show that the agent was able to use this method to autonomously learn to perform the tasks.

A Goal-Directed Visual Perception System Using Object-Based Top–Down Attention

Yuanlong Yu; Mann, G.K.I.; Gosine, R.G. Page(s): 87 - 103 (PDF)
Abstract: The selective attention mechanism is employed by humans and primates to realize a truly intelligent perception system, which has the cognitive capability of learning and thinking about how to perceive the environment autonomously. The attention mechanism involves the top-down and bottom-up ways that correspond to the goal-directed and automatic perceptual behaviors, respectively. Rather than considering the automatic perception, this paper presents an artificial system of the goal-directed visual perception by using the object-based top-down visual attention mechanism. This cognitive system can guide the perception to an object of interest according to the current task, context and learned knowledge. It consists of three successive stages: preattentive processing, top-down attentional selection and post-attentive perception. The preattentive processing stage divides the input scene into homogeneous proto-objects, one of which is then selected by the top-down attention and finally sent to the post-attentive perception stage for high-level analysis. Experimental results of target detection in the cluttered environments are shown to validate this system.

Editor : Pierre-Yves Oudeyer, INRIA Editorial Assistants : Adrien Baranes and Paul Spears
The AMD Newsletter is financially supported by the IEEE CIS and INRIA