IEEE CIS
AMD Newsletter
Robotics
Psychology
Neuroscience
Machine Intelligence
The Newsletter of the Autonomous Mental Development Technical Committee

VOLUME 10, NUMBER 1 ISSN 1550-1914 Spring 2013

Editorial

In socially assistive robotics, a field of immense societal value, robots need to adapt continuously to the particularities, preferences and evolving needs of the human(s) they were made to help, as well as to cluttered human environments. These requirements meet the objectives of developmental robotics, where methods for lifelong incremental learning are studied, with a strong guiding role for developmental constraints coming both from within and outside the learner. Across the dialog initiated by John Weng, at the crossroads of assistive and developmental robotics, and with contributions from Yiannis Demiris, Adriana Tapus, Manuel Lopes, Katharina Rohlfing and Britta Wrede, Anthony Morse and Yoonsuck Choe, everyone converges to this general vision. Yet, many important and complicated issues remain open. In particular, while there are arguments stating that concepts and representations in assistive robots should be emergent and not hand tuned by engineers, it is also crucial, in such application fields, that safety, trust, acceptability and usefulness for the humans can be ensured. While in a distant future we can see how these two lines could be converging, in the short term they can be incompatible for real-world applications. The dialog in this issue of the newsletter provides enlightening visions about this challenge.

Then, Peter Ford Dominey initiates a new dialog over a crucial topic for which much remains to be discovered and understood: "How are grammatical constructions linked to embodied meaning representations?". In particular, what are the developmental processes leading to the acquisition of grammatical construction of increasing complexity, and to their associated mental simulation capabilities? How is it possible to take a computational modeling approach to these questions?

Those of you interested in reacting to this dialog initiation are welcome to submit a response by September 15th, 2013. The length of each response must be between 600 and 800 words including references. Contact: contact pierre-yves.oudeyer@inria.fr

— Pierre-Yves Oudeyer, Inria, Editor

Message from the Chair of AMD Technical Committee

We are all looking forward to participating in the joint ICDL-EpiRob Conference, which will be held in Osaka on 18-22 August 2013. This year we will have keynote presentations by Anne Fernald (Stanford University, USA), Herbert Jaeger (Jacobs University, Germany), Yasuo Kuniyoshi (The University of Tokyo, Japan) and Ichiro Tsuda (Hokkaido University, Japan). The conference will also have satellite events: the final conference of the European doctoral training network in developmental robotics (www.robotdoc.org) held on August 16-18 and the 5th Symposium on Cognitive Neuroscience Robotics (http://www.gcoe-cnr.osaka-u.ac.jp/5th-symposium/en/) on August 22.

In Osaka, on August 21, we will also have the Annual Meeting of the IEEE AMD Technical Committee, which this year will be held together with the ICDL Governing Board and the EpiRob Steering Committee. This will be an opportunity to discuss the establishment of a new Governing Body for the joint ICDL-EpiRob community. The AMD Technical Committee is also considering changing the "Autonomous Mental Development" name, to choose a more general and self-explanatory name. The new names could apply to all our activities, i.e. the Technical Committee, the IEEE TAMD journal and the joint ICDL-EpiRob conference.

The new, unifying name will increase the visibility of our work beyond the current community, including all disciplines interested in development and learning. This name change will also be part of the discussion during the joint meeting on August 21.

Finally, plans have already started to host the 2014 ICDL-EpiRob conference in Genoa, Italy, with the organisation team led by Giorgio Metta and Giulio Sandini. An update on the 2014 event will be communicated in Osaka.

— Angelo Cangelosi, chair of AMD TC

Dialog Column

Modeling AMD and its Application to Assistive Robotics: Closed Skull or Not?

Juyang Weng
Dept. of Computer Science and Engineering
Michigan State University, USA

Task-nonspecific developmental programs for Autonomous Mental Development (AMD) have been raised for over a decade [1]. The term "autonomous" in AMD was insisted on by the late Developmental Psychologist Esther Thelen. She said to me that without autonomy the human development program (DP) could be mistreated as simply rolling out functions and behaviors. For example, autonomous actions play an important role in explaining the A-not-B error [2]. However, there is no consensus yet about what the term "autonomy" means exactly.

What do we mean by autonomy in AMD? One may say that a developmental agent must be autonomous throughout development. However, this is not very clear. For example, suppose that a teacher shows to a child how to draw a house by holding and guiding the movement of his hand that holds a pen. In developmental psychology, it is called passive learning. According to the refined eight types of learning [3], it is called type 2 learning, i.e., emergent-imposed-communicative learning. Is this passive learning consistent with the term "autonomy" that Esther Thelen insisted on? Note that the effector is not autonomous in this case.

Therefore, I proposed that "autonomous development" means "development with a closed skull": it is not allowed for a human teacher or programmer to open the learner's skull, e.g., supervise the connections and responses inside the "brain" while the agent is learning. This autonomy still allows and encourages human teachers to manually interact with the brain, but only through the brain's sensory port and the effector port [3]. The mechanisms allowing an autonomous robot to adapt continuously to the evolving skull-closed social interaction with humans have been argued to be key in establishing truly natural social interaction and mutual understanding between humans and robots [4]. Yet, some have also argued that autonomous developmental approaches may sometimes be problematic for socially assistive robots [5]. Is skull-closed development inappropriate for socially assistive robots that serve vulnerable users? What do we mean by an agent being brittle? Can skull-open machine learning resolve the high brittleness? Does open-skull interaction provide a more successful route toward robust and safer assistive robots for vulnerable users? How do we understand the fact that human nurses use closed-skull development?

References:

[1] Weng, J., McClelland, J., Pentland, A., Sporns, O., Stockman, I., Sur, M., & Thelen, E. Autonomous Mental Development by Robots and Animals, *Science*, 291 (5504), pp. 599-600, 2001.
[2] Smith, L. B., Thelen, E., Titzer, R. & D. McLin. Knowing in the context of acting: The task dynamics of the A-not-B error, *Psychological Review*, 106, pp. 235-260, 1999.
[3] Weng, J. *Natural and Artificial Intelligence: Introduction to Computational Brain-Mind*, BMI Press, Okemos, Michigan, 2012.
[4] Rohlfing, K. & Wrede, B. What Novel Scientific and Technological Questions does Developmental Robotics Bring to HRI? Are we Ready for a Loop? Dialog column of *IEEE CIS AMD Newsletter*, 9(1), pp. 3-5, 2012.
[5] Tapus, A. Is Developmental Robotics a Solution for Socially Assistive Robotics? Dialog column of *IEEE CIS AMD Newsletter*, 9(1), pp. 9-10, 2012.

The Altruistic Wheelchair: Development Meets Assistive Robotics

Yiannis Demiris
Personal Robotics Laboratory, Imperial College London

In developmental robotics we design mechanisms that enable robots to explore themselves and their environment, and exploit the acquired knowledge to advance criteria considered important within this environment. This knowledge acquisition is frequently facilitated through social mechanisms, for example imitation learning, using methodologies inspired by child development [1]. The intersection of developmental robotics with assistive robotics highlights two further interesting challenges: altruism and transparency.

Firstly, in assistive contexts the developmental criteria considered important are *allocentric*, focused on the developmental needs of the assisted human, while the egocentric ones (the ones of the robot) are necessarily taking a secondary role. This is particularly important in lifelong scenarios where the human is heavily dependent on the robot assistance for an extended period of time. For example, in our own work, insights from developmental psychology have been helpful in designing robotic paediatric wheelchairs that can assist and encourage very young disabled children to safely explore their environments [2], in order to avoid developing learned helplessness. While robot-initiated environment exploration would be beneficial for the robot's development (for example the acquisition of more detailed environmental information), the robot will need to take into consideration the child's short-term and lifelong developmental needs (for example by inferring the child's zone of proximal development [3]) and thus promote child-led exploration.

Secondly, our experience deploying robotic assistants in paediatric hospitals (both in the context of the paediatric robotic wheelchairs above, as well as in the context of robotic long-term companion humanoid robots [5]) have taught us that the involvement of sensitive human populations necessitates more *transparent* action selection mechanisms so that validation of these mechanisms and potential intervention by experienced human personnel is facilitated. The exact mechanism for this second, more pragmatic aspect, can vary since transparency can be realized either through the externalization and modification of the acquired knowledge through natural interfaces (for example speech, or visually-perceived task demonstrations), in a manner similar to John Weng's example of human nurses, or via directly observable and potentially modifiable internal representations. But in either case, the presence of such mechanisms greatly facilitates adoption of developmental assistive robots in real environments for the benefit of sensitive populations.

Ultimately, the assistive robotics field provides an intriguing twist to the design requirements of our developmental mechanisms. Viewing the human and the robot as a developmental system, we can control the behavior of one part of the system (the robot), but the main metrics for measuring success in assistive contexts are human-centered [4], the part of the system you cannot directly control. This altruism must be transparent and verifiable if we target to deploy these developmental systems in real-world assistive contexts.

References:

[1] Demiris, Y. and Meltzoff, A. The Robot in the Crib: A developmental analysis of imitation skills in infants and robots. *Infant and Child Development*, 17, pp. 43-53, 2008.
[2] Soh, H. and Demiris, Y. Towards Early Mobility Independence: An Intelligent Paediatric Wheelchair with Case Studies, *IEEE/RSJ IROS2012 Workshop on Navigation and Manipulation Assistance for Robotic Wheelchairs*, Oct 2012.
[3] Demiris, Y. Knowing when to assist: developmental issues in lifelong assistive robotics, in *Proceedings of the 31st Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC 2009)*, pp. 3357-3360, Minneapolis, Minnesota, USA, September 2-6, 2009.
[4] Carlson, T and Demiris, Y. Collaborative Control of a Robotic Wheelchair: Evaluation of Performance, Attention and Workload, *IEEE Transactions on Systems, Man and Cybernetics B*, 42(3), pp. 876-888, 2012.
[5] Ros, R., Baroni, I., Nalin, M. and Demiris, Y. Adapting Robot Behavior to User's Capabilities: a Dance Instruction Study, *ACM/IEEE Human-Robot Interaction conference (HRI-2011)*, Lausanne, pp. 235-236, 2011. Part of the EU FP7 ALIZ-E project, www.aliz-e.org

Can We Trust Closed-Skull Socially Assistive Robots?

Adriana Tapus
Robotics and Computer Vision Lab ENSTA-ParisTech, Paris, France

Robots are more and more present in our daily life; they have to move into human-centered environments, to interact with humans, to obey some social rules so as to produce an appropriate social behavior in accordance with human's profile (i.e., personality, state of mood, and preferences), and to provide assistance to vulnerable users. More and more research works are trying to step up from short-term and task specific approaches towards the development of emerging solutions for life-long and adaptive interactions [1], [2], [3].

First of all, one of the requests of Prof. Weng was the definition of "open skull" and "closed skull". The idea behind "closed skull" is how to continuously and robustly learn from almost nothing a variety of unconstrained tasks and reliably adapt to unseen contexts similarly to the way humans do it. With this in mind, I understand "closed skull" as an autonomous life-long learning and decision process with no human intervention (i.e., continuously learn and decide what to learn, when to learn, and how to behave appropriately). On the other hand, I see an "open skull" system as a system that can autonomously act, learn, and interact with human peers for a variety of specific tasks. However, if the context for which the system was tailored is highly dynamic, it is allowed for the engineer to freely "open the skull at night" to study and improve the system.

As previously stated in my response [4] to the point raised by Rohlfing and Wrede in [5], I agree that the robot has to be capable of life-long incremental learning in order to better adjust its behaviors and provide an appropriate response as a function of various situations and tasks. Humans and environmental factors can be unpredictable and of course the brittle effect can appear. A system cannot be built and handcrafted to be perfect and this is an aspect that I totally share with Prof. Weng. Nevertheless, I believe that a system that can evolve with time and incrementally learn new inputs can work perfectly. In his research works [6], Brooks advocates that a complex and sophisticated intelligent robotic system can be developed by the incremental addition of individual layers of situation-specific control systems. Or maybe as suggested in [7], develop a solution for perturbation-tolerant autonomous system whereby the robot can infer whether or not it is achieving its goals, and if not, trying a potential "Plan B", or more or less random variations in behaviour, and ask for help, or use trial-and-error.

Assistive robots have to provide assistance to vulnerable users in a context-specific scenario, and need to respond to both short-term changes that represent individual differences and long-term changes that allow the interaction to continue to be engaging over a period of months and even years. By using a "closed-skull" development, the system must know robustly when, who, and how to assist the individuals [2]. I believe that machines' perception still lacks reliability and robustness in complex natural environments and that's a bottleneck for performing tasks that humans do effortlessly. So, at the question raised by Weng: "Is skull-closed development inappropriate for socially assistive robots that serve vulnerable users?", my concern is about how are we able to measure what the skull-closed robot has learned and its evolution in time and how can we guarantee an appropriate safe robot behavior that will serve the users' needs. One fundamental concern is whether such robots can be trusted (trustworthiness quantified in terms of usefulness, safety, and predictability).

The tasks and the trainings required by the therapists/nurses are very different from one medical condition to another and from one user to another, making the interaction even harder and challenging. While the validation domains and thus some user populations (e.g., the elderly, post-stroke patients, children with autism) are quite diverse, the principal facets of the human-robot interaction needed for Socially Assistive Robotics (SAR) care are quite similar. The system complexity has to be variable within a wide range of individual needs. The interactions require adaptation over time so as to enable training behavior (cognitive or physical) toward long-term performance goals. With this view in mind, I believe "open-skull" interactions provide a sufficient and yet more successful route toward robust and safer assistive robots for vulnerable users.

References:

[1] Baxter, P., Belpaeme, T., Canamero, L., Cosi, P., Demiris, Y. and Enescu, V. Long-term human-robot interaction with young users, in *Proceedings of the ACM/IEEE Human-Robot Interaction conference (HRI-2011) (Robots with Children Workshop)*, Lausanne, Switzerland, 2011.

[2] Demiris, Y. Knowing when to assist: Developmental issues in lifelong assistive robotics, in P*roceedings of the Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*, Minneapolis, MN, pp. 3357–3360, 2009.

[3] Tapus, A., Tapus, C. and Mataric, M.J. Long term learning and online robot behavior adaptation for individuals with 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, 62, pp. 389–401, 2010.

[4] Tapus, A. Is developmental robotics a solution for socially assistive robotics? *Dialog column of IEEE CIS AMD Newsletter*, 9(1), pp. 9–10, 2012.

[5] Rohlfing K. and Wrede,B. What novel scientific and technological questions does developmental robotics bring to hri? are we ready for a loop? *Dialog column of IEEE CIS AMD Newsletter*, 9(1), pp. 3–5, 2012.

[6] Brooks, R. A robust layered control system for a mobile robot, *IEEE Journal of Robotics and Automation*, 2(1), pp. 14–23, 1986.

[7] Anderson, M. L. and Perlis, D. R. Logic, self-awareness and self improvement: The metacognitive loop and the problem of brittleness, *Journal of Logic and Computation*, 14, 2004.

Assistive Robotics: Domain Knowledge, Experience and Interaction with the User

Manuel Lopes
Flowers project-team, Inria and Ensta ParisTech, France

In all the recent successes of deployment of robots, we can see that the state of the art is still single purpose model based machines. Industrial robots, drones, tele-operated surgical robots, all correspond to domains we can model accurately and where we can formally understand the task. Can we develop an assistive robot? Do we know how to deal with vulnerable users? Are they all equal or do we need personalization? Does providing a behavior that is good for the ‘’average’’ user better than nothing, or will system without personalization just break apart? The inexistence of assistive robotics in the real world, besides research studies, is probably a proof that a model based approach is not able to provide a solution.

There are a few problems where learning and/or development seems to be the only option: highly non-linear systems (e.g. discontinuous, turbulent) and when the human factor is very strong. In assistive robotics, each particular user requires a particular adaptation, with specific representations: what is good for a person might be bad for another. Such knowledge can only be acquired by interacting with each particular user. Even the measures of success are very unclear, different people express joy, irony or sarcasm in different ways, and even language is often ambiguous. A robust system cannot assume that any of its representations and parameters is fixed - we are no longer in a factory where the operator must adapt to the robot. For this reason a perspective that considers personalization through the adaptation of interaction protocols, representations and preferences, will provide great improvements in such systems.

We, and other researchers, have been developing methods for adaptive learning of interaction protocols [1], optimal teaching [2] and interactive scenarios [3,6] to ensure that each system is particularly adapted to a single user. It is already clear that each person behaves differently during interactions [4] and that without personalization the results can be worse for everyone that a truly adapted system [5].

In this dialog, the discussion focuses on which approach – open skull versus closed skull – will be best for this problem. This distinction is not clear and just follows the same lines as other discussions [7,8]. In terms of understanding the mechanisms of a machine (or of a human) we need both. In biological systems we use EEG, TMS, single cell recording, drugs and psychology. From a computational point of view, control theory tells us that some effects cannot be observed from the outputs, and some states cannot be directly controlled from the inputs. It should be clear now that learning from nothing is impossible and trying to model everything is also impossible.

The challenging problem of assistive robotics will need a truly developmental approach. Domain knowledge will be integrated in the system from experts and prototypical interactions will be designed. But, due to a lack of full knowledge of the domain, the system will develop further capabilities mostly autonomously and through social interaction with each individual user. At the same time, an expert human is required to be on-the-loop (probably not necessarily in-the-loop) to improve safety and allow any kind of intervention (directly on the machine or through its inputs) when the system goes away from what the expert therapist thinks is appropriate [3].

References:

[1] Lopes, M., Cederborg, T. and Oudeyer, P-Y. Simultaneous Acquisition of Task and Feedback Models, *IEEE - International Conference on Development and Learning (ICDL)*, Germany, 2011.
[2] Cakmak, M. and Lopes, M. Algorithmic and Human Teaching of Sequential Decision Tasks, *AAAI Conference on Artificial Intelligence (AAAI)*, Toronto, Canada, 2012.
[3] Mason, M. and Lopes, M. Robot Self-Initiative and Personalization by Learning through Repeated Interactions, *6th ACM/IEEE International Conference on Human-Robot (HRI'11)*, Lausanne, Switzerland, 2011.
[4] Knox, B., Glass, B.D., Love, B.C. Maddox, W. T., Stone, P. How Humans Teach Agents, *International Journal of Social Robotics*, 4(4), pp. 409-421, November 2012.
[5] Lee, J. and Brunskill, E. The Impact on Individualizing Student Models on Necessary Practice Opportunities, *International Conference on Educational Data Mining* (EDM) 2012.
[6] Thomaz, A., Breazeal, C. Experiments in socially guided exploration: lessons learned in building robots that learn with and without human teachers, *Connection Science* 20 (2-3), 91-110, 2008.
[7] Breiman, L. Statistical Modeling: The Two Cultures, *Statist. Sci*. 16(3), pp. 199-231, 2001.
[8] Elman et al.. *Rethinking Innateness: A Connectionist Perspective on Development*, MIT Press, Cambridge, MA, 1996.

Recurring Interaction Patterns Enable the Emergence of Concepts from Previous Interaction: the Importance of Long-Term Memory in Interactive Learning


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

In the AMD Newsletter Fall 2012, Weng argues that current implementations with the goal of leveraging pragmatic information from human-robot interaction for the learning of concepts are still based on symbolic representations. We admit that in current approaches, interaction relies on symbolic concepts. In effect, they still lack by far the smoothness with which humans are able to adapt to and learn from previous interactions. However, the approaches are helpful to demonstrate the power of interactional cues, and how they can perpetuate learning processes.

In our current approaches [1], we dissolve the necessity to supervise the brain by proposing "frames" in developmental robotics. Frames are recurring interaction patterns (e.g. [2]) and have been observed to play an important role in parent-infant interactions for the emergence of concepts from interaction. An example is labelling an object: A 'labelling-frame' consists of components such as calling somebody's name or looking at the person, looking at and pointing to an object and labelling it. Thus, once such a sequential structure is established, the only new element within it is the new label for a particular object. In this sense, embedding a new word within a familiar frame results in the reduction of the information load on the child. We think that much more fine-tuned frames can be observed and are often developed in idiosyncratic interactions first. They allow, for example, detailed corrections of infant's imitation attempts – helping the infant to understand underlying semantic concepts such as goals, constraints or means.

In this vein, results on low-level visual analyses of parent-child interactions indicate that the application of fundamental biases, such as the focus on synchrony or contingency as well as visual saliency, may allow to bootstrap the first frame-like interaction structures: E.g., [3] have shown how at the end of an action presentation, parents are suppressing their own movements – with the

effect that the visually most salient region moves from the parent's hands and face to the created object situation. This frame, thus, emphasises the achieved goal situation without making use of symbolic representations.
Careful research is needed to better understand how such frames emerge in interactions and how they can be learned with existing machine learning approaches. In [4], we show how mothers educate and reward the child's eye behaviour to become a communicative means in form of an eye contact.

More complex frames are reported e.g. in [5] where fundamental biases such as synchrony between action and speech achieve a segmentation of the demonstrated action stream. More detailed analyses will be needed to investigate how such attentional processes towards binding speech and action emerge and how interactions shape the concrete parameterisation of this process. For example, the proposed approach makes use of the dominant role of speech which may emerge because of its transient nature: while visual results remain perceivable after an action has been carried out, speech does not.

We propose that a range of such fundamental biases exist (e.g. sensitivity towards biological motion, preference for speech) that help to shape the first frames in interaction with a caretaker. However, the recurrent interaction is the essence of the frames. In the repeated interaction, the child will look for the invariant aspects and her or his successful exploration crucially depends on the sensitivity of the caregiver to the infant's reaction to patterns of interaction. Crucially, they not only attract the child's attention in an effective way but also evoke an interactional expectation [6].

Frames, thus, are co-shaped entities emerging first between an infant and an individual caretaker as interaction protocols. The transfer to different individuals may be modulated by abstractions or conventionalisation of the learned concepts, e.g. the concept of a „face" or a „voice", which have symbolic quality [4].

While computational analyses have shown that at least some such frames can be detected automatically, it is still unclear how the concept of frames can be leveraged for robotic learning to embed standard learning algorithms in such adaptive frame-like structures. Frames are multi-modal recurring interaction patterns that can encompass an immense range of modalities, structures and functions (e.g. shifted saliency due to frozen movements at the end of an action, attention shift to acoustic packages due to synchrony of speech and action events, reward for eye gaze behaviour through interactive cues such as smiles or vocalisations etc.) that operate through the fine-grained manipulation of (and reward for) the learner's attention. To learn them, the robot must be able to engage in a recurrent social interaction. The learning mechanism itself needs to be able to model temporal structure of interaction, to take rewards (with a certain time delay) into account and to allow for incremental growth by e.g. learning new rewards (e.g. smiling, eye gaze, face of mother/father, voice of mother/father) from basic reward mechanisms (e.g. contingency).

The shortcoming of current learning approaches are due to the fact that learning is bound to one pre-defined interaction protocol and tuned to the underlying learning algorithm. In contrast, we propose that by enabling the learning system to learn (idiosyncratic) interaction frames through recurrent interactions, not only the teaching situation itself can be made less brittle but, more importantly, that a qualitatively new level of learning and understanding can be achieved.

References:

[1] Wrede, B., Rohlfing, K. J., Steil, J. J., Wrede, S., Oudeyer, P.-Y. & Tani, Y. Towards robots with teleological action and language understanding, in *Proceedings of the IEEE-RAS International Conference on Humanoid Robots (HUMANOIDS)*, Osaka, 2012.
[2] Bruner, J. S. *Child's talk: Learning to use language New York*, Norton, 1983.
[3] Nagai, Y. & Rohlfing, K. J. Computational analysis of Motionese toward scaffolding robot action learning, *IEEE Transactions on Autonomous Mental Development,* 1, pp. 44–54, 2009.
[4] Nomikou, I, Rohlfing, K. J., Szufnarowska, J. Educating attention: recruiting, maintaining, and framing eye contact in early natural mother-infant interactions, *Interaction Studies*, 2013.
[5] Schillingmann, L., Wrede, B. & Rohlfing, K. J. A computational model of acoustic packaging, *IEEE Transactions on Autonomous Mental Development*, 1, pp. 226–237, 2009.
[6] Gergely, G. & Csibra, G. Sylvia's recipe: The role of imitation and pedagogy in the transmission of cultural knowledge. *Roots of human sociality. Culture, cognition and interaction,* Enfield, N. J. & Levinson, S. C. (Eds.), Oxford / New York, Berg, pp. 229–255, 2006.

Thinking Outside the Skull...

Anthony F. Morse
Cognition Institute
Centre for Robotics and Neural Systems
Plymouth University

In the quest for autonomy a myriad of methods and approaches have been proposed and refined and more often than not spun off in various directions spawning technologies and industries. Yet in truth, artificial autonomy in all but the simplest of toy problems remains elusive. While I wholeheartedly endorse the proposal that autonomy in development must be closed skull, with respect to any particular application it not always so clear that; a) autonomy must be all or nothing, and b) that autonomy is always desirable.

So why then these concerns? Yes, autonomy is crucial but perhaps the worry is one of design bias; by interfering with the internal dynamics and specifying 'labels' we risk imposing our assumptions, our naive reductions, and our naive methods on the task to be accomplished. Without doubt these impositions narrow the niche but they can equally advance abilities significantly within that narrow niche. For example, as humans we do not use an exhaustive heuristic search when we play chess, but if you want a machine to play chess then this is not a bad choice of method, it just won't be much good at anything else. Assistive robotics is a prime and very timely example in which huge research resources are pilling in and producing genuinely useful robots, many of which are already in use. Examples include robots helping nurses lift heavy patients, and even robots that can allow paralysed people to walk [1]. This is yet another spin off, productive, profitable, and beneficial, but ultimately a path diverging (for now) from autonomy. But it is not clear that we should object on grounds of lacking autonomy.

I share the concern that open skull engineering may produce brittle results, but I am equally concerned by the relative lack of function of more closed methods. Most closed methods simply lack the scalability to escape toy world problems. The issue of design bias is also a matter of degree, even in closed methods design bias can still be observed, in pre-shaping the structures that serve a hierarchy, in selection of the parameters of associative learning, even so simply in the choices of pre-processing data, and choosing sensory and motor systems. As a community we must be aware of these issues.

Our understanding of embodiment is also transforming how we think about cognition; cognition is no-longer in the skull but is extended [2], trading computation with the body and environment, briefly coupling with natural and artificial technologies to form new dynamic systems. In this way, even radically open design in the skull may ultimately leave much of the cognition and autonomy outside, beyond that with which we tinker.

In conclusion I am torn, I want to see a world of truly useful cognitive and autonomous robots but that vision seems far away. On the immediate horizon is a world of truly useful robots but I fear their inevitable success will distract from the methods leading to further advances. I seek to understand development and so avoid both implausibly supervised methods and those that do not scale [3], but we are an interdisciplinary community, and with that we must recognise that we raise different questions and seek different goals. Questions such as "how can we make robots that help people?" and "how can we make robots that autonomously mentally develop?" are clearly different. As researchers we must select our methods carefully so as not to undermine the questions we are asking. Just as planes can fly without flapping their wings, robots can and (without exception) do function in radically different ways from our own cognition and autonomy.

References:

[1] Esquenazi, A., Talaty, M., Packel, A., & Saulino, M. The ReWalk Powered Exoskeleton to Restore Ambulatory Function to Individuals with Thoracic-Level Motor-Complete Spinal Cord Injury, *American Journal of Physical Medicine & Rehabilitation*, 2012.
[2] Clark, A., and Chalmers, D. The extended mind, *Analysis*, pp. 7-19, 1998.
[3] Morse, A. F., de Greeff, J., Belpeame, T., & Cangelosi, A. Epigenetic robotics architecture (ERA), *Autonomous Mental Development, IEEE Transactions on*, 2(4), pp. 325-339, 2010.

How Will the Emergent Representations Be Used, From Within the Skull?

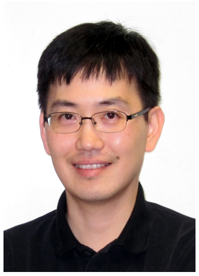

Yoonsuck Choe
Department of Computer Science and Engineering
Texas A&M University, College Station, TX, USA

Forming emergent representations in a closed skull is one thing; learning how to use those representations is another. In the dialog initiation, Weng only addresses the former problem. I would like to argue that the latter is as important as the former, and considering the latter will help answer the many questions Weng posed in his dialog initiation.

First, is the "closed skull" approach better than the "open skull" approach? Yes, it is, but not because of the obvious reason that closed skull is how it is done in nature. I believe the closed skull approach is very important because taking that perspective helps us identify the real, immediate problems faced by the skull-enclosed brain and their potential solutions – not the artificial problems and solutions we, the scientists, usually come up with. For example, consider neural spike analysis. Scientists pose the problem as an encoding-decoding problem (cf. "grounding" [1]). Thus, the solution involves the simultaneous observation of the external stimulus and the internal spikes: reverse correlation (e.g., spike-triggered average). However, the solution is unrealistic in view of a closed skull, since the external portion of data needed for the solution (the external stimulus) is not directly available from inside the skull. Thus, this solution, although highly accurate, cannot be employed by the brain itself. Now, if we take the closed skull perspective, initially, spike decoding appears to be an impossible problem to solve: How can the brain know what the somewhat arbitrary spikes mean, without any external reference? It turns out that there is an easy solution, and it involves the motor system. By moving and observing the systematic changes in the internal spikes, the brain can understand how stimulus properties encoded in the spikes correlate to specific kinds of motor behavior, all without peeking outside of the skull (see [2],[3] for details). Coming up with this kind of solution is impossible without taking the "closed skull" perspective, thus justifying my initial assertion above. Also, such a motor action-based solution provides an answer to the question of " how to use those representations".

Next, Weng asks if skull-closed development is more appropriate for assistive robots and poses related questions regarding the brittleness of representations and robustness of assistive robots using such representations. It is conceivable that once scientists and engineers find out how the brain works and through closed-skull development train their first AMD-based robot, we can simply generate copies of the robot in an open-skull manner. These open-skull robot copies can be robust, however, this assumes that AMD-based robots can stop their development process at a certain time (at which point the copies are made) and they will remain robust thereafter. However, I don't think this will be the case. A truly robust AMD-based robot should continually develop its emergent internal representations, and this seems to be the only way to avoid brittleness, where we can think of brittleness as the inability to cope with unforeseen situations. Because the robot's internal representations will be continually changing dependent on the experience of the individual robot, the robot must maintain its closed-skull development so that the meaning of the internal representations can be figured out by the robot itself through continual motor exploration. Humans can intervene in an open-skull manner, but the task will be very tedious and laborious, so it would be much better to implant a sensorimotor learning mechanism so that the robot can figure out by itself how to use the internal emergent representations.

References:

[1] Harnad, S. The symbol grounding problem. *Physica D*, 42, pp. 335-346, 1990.
[2] Choe, Y. Action-based autonomous grounding. In Modayil, J., Precup, D., and Singh, S., editors, *AAAI-11 Workshop on Lifelong Learning from Sensorimotor Experience*. Palo Alto, CA: AAAI Press. AAAI Workshop Technical Report WS, pp. 11-15, 2011.
[3] Choe, Y., Yang, H.-F., and Eng, D. C.-Y. Autonomous learning of the semantics of internal sensory states based on motor exploration. *International Journal of Humanoid Robotics*, 4, pp. 211-243, 2007.

Reply and Summary
Modeling AMD and Its Application to Assistive Robotics: Closed Skull or Not?

Juyang Weng
Department of Computer Science and Engineering,
Cognitive Science Program, Neuroscience Program,
Michigan State University, East Lansing, MI 48824 USA

At the current state of awareness, I raised this issue but got little support in this dialogue. Yoonsuck Choe is an exception. Choe wrote "forming emergent representations in a closed skull is one thing; learning how to use those representations is another." The former makes the latter possible. With a closed skull through life-long AMD, the internal representations must satisfy the following conditions: They not only autonomously emerge (which traditional neural networks can claim to do) but also satisfy many conditions that traditional neural networks cannot do well with a closed skull, such as modality independence (vision, audition, text), task-independence (from simple to complex), environmental openness (directly learn from cluttered environments), top-down attention (e.g., attend objects in cluttered settings based on emergent intent), concept abstraction (e.g., concept as states from varied stimuli), concept generalization (e.g., error-free for state-equivalence sequences that the agent has not learned), bootstrapping (reuse the same experience/data to improve), scaffolding (use earlier learned skills for learning later skills), and transfer (skills learned in one setting are automatically used properly for other different settings). Choe was correct in raising the representation emergence issue, but an open skull (as many traditional neural networks have done) has led to many aspects of representations that should not be handcrafted in the first place (e.g., impose that the function of a module is to perform edge detection), since such a practice leads to a highly brittle system.

I like Choe's example of encoding-decoding problem. Indeed, with closed-skull under consideration, we can see that encoding-decoding is not what a brain does at all. Indeed, concurrent actions give class labels for perception in many cases. However, I proposed further in [3] that the motor area in the brain does not only decoding from encoded sensory information (i.e., not feedforward from sensor to motor). But rather, internal representations as "bridges" that take sensory inputs and motor inputs as two source banks (not one) to catch statistics of many different sensory receptive fields (one bank), effector receptive fields (another bank), and lateral receptive fields (in the bridge). Furthermore, the emerging bridge representations are used not only by the motor "bank" but also the sensory "bank" for their own real-time predictions. Choi was correct in stating that any static copy of a brain will not be robust either. I also agree with Choi in that any human intervention is tedious and laborious. But I also argue that any human intervention is damaging since it is grossly non-optimal into what was already optimal. According to the DN theory [3], with its closed skull the brain's emergent internal representations are almost optimal.

Thus, in response to the points raised by Adriana Tapus, I would argue that in HCI it appears to be beneficial, at least in the future when developmental robots have become a common reality, not to allow the programmer of a developmental robot to "open the brain skull at night" although his motivation is to "improve". If the purpose is to produce a visualization of the brain representations for an improvement of future developmental programs (DP) of other robots, I am for it because its DP is handcrafted. However, if the purpose is to tune some parameters of this brain's internal representations, then I have reservation because the brain representations are not handcrafted - too complex to handcraft (or hand-tune) well. Let's imagine that a brain surgeon wants to perform brain surgery on a nurse because she does not do a satisfactory job in taking care of elderly people. The surgeon would likely produce more safety hazard if we allow him to do this. Similarly, I think that socially assistive robots would be more useful and safer to elderly if they can learn directly in their assistive working environments like how a nurse learns during her job training.

Katharina Rohlfing and Britta Wrede's "frames" seem to be still brittle since their elements are based on human handcrafted concepts. There are many problems in such symbolic representations [4]. For example, they can result in high brittleness due to the intrinsic exponential complexity of symbolic representations. Consider that a nurse's brain is skull closed. For instance, how does she learn concepts like "silent" and "non-silent" which is handcrafted into the used speech recognition system? I argue that such concepts are not determined by a static threshold that is innate to the brain. Every concept depends on context. For example, a population that extols its autocratic king is silent, but a softly ticking mechanic watch is not silent. In DN, to fully take into account contexts, every action emerges (motor-supervised learning is a type of emergence) in actual context, not in terms of a set of handcrafted rules. Yet, Rohlfing and Wrede insist on one very important issue that has indeed to be solved: interaction protocols should not be pre-defined for any AMD system. The BMI 831 course in the 2013 summer school of the Brain-Mind Institute provide detailed studies about how animals interactively learn communication protocols.

I agree with Manuel Lopes in that "a model based approach is not able to provide a solution". AMD is regulated by the complex genome program (DP). I do not agree on the necessity that "domain knowledge will be integrated in the system from experts and prototypical interactions will be designed". Only some innate behaviors (like sucking and rooting in the case of human newborns [1]) that may assist early autonomous development may be needed. The primary reason is that handcrafted domain knowledge and prototypical interactions are damaging for an AMD system because they are not only suboptimal but also highly brittle. For a high robustness, expert knowledge should be taught from the brain's external environment through brain's sensors and effectors, instead of being implanted directly into a skull-open brain. The former is desirable since the system is subject to realistic operational environments during the process of learning and examination but the latter does not necessarily require that.

Anthony Morse wrote: "without doubt these impositions narrow the niche but they can equally advance abilities significantly within that narrow niche". Yet, I argue that under an optimal developmental paradigm, manual imposition not only damages the optimality but also introduces a high brittleness. As a clean problem, the example of playing chess does not apply to the development of brain functions since assistive robots must deal with highly muddy problems. For clean problems we do not need AMD. His statement "most closed methods simply lack the scalability to escape toy world problems" is incorrect. For example, the closed-skull developmental algorithm DN [3] learned directly from cluttered environments to solve "where" and "what" problem conjunctively for multiple general objects. By this capability alone, the term "toy" does not apply since a cluttered environment contains many unknown objects. When a baby visually learns a toy in a cluttered natural environment, its skull-closed brain is not dealing with a "toy problem" because of the background involved among many other challenges. A survey on existing computer vision methods (e.g., handcrafted graphic models) can tell us that there is no open-skull method that can effectively learn general objects (non-face) directly from cluttered environments. This means that all existing handcrafted computer vision methods (yes, skull open) are intrinsically not scalable because of the brittleness of human handcrafting. There exists no truly general-purpose vision algorithm yet, skull-open or skull-close. Although the DN seems to be the first beginning one, we have still a lot of work to do before we can have a truly general-purpose vision algorithm. For example, emergence of concept hierarchy needs life-long development and a brain-like computational resource before this highly desirable capability can be convincingly demonstrated. Even a fruit fly has a huge resource compared to its receptors. Although we can never faithfully model a brain, I feel that a practical general vision problem cannot be solved by a system that does not have a brain-like amount of computational resource. The brain devotes about a half of the cortical area for visual information processing [2] and I feel that a normal brain is nearly optimal for its species environment and experience. Our experimental results used a very limited amount of computational resource that is far from that of the human brain, in the current necessary stage for the proof of concept. The DN approach appears to be the only exiting one that is scalable for general-purpose vision. It is also under a framework of optimality. But it is unknown at this time how soon such an approach can reach the visual performance of an animal (e.g., cat).

Yiannis Demiris raised two issues that I think are indirectly related to the issue of closed skull: allocentricity and transparency. According to developmental psychology, egocentricity develops earlier than altruism (or earlier than a capability to think in the position of others). As I argued above, for clean tasks we probably do not need AMD. Here, we are particularly interested in muddy tasks that traditional assistive robots do not perform well. I guess that the development of sense of self is necessary for extending to a sense of others, since self is a special case of multiple agents in the environment. I here raise a question: Can a robot be sufficiently smart if it was programmed to be truly unselfish? From our current studies of reinforcement learning with DN, I guess that selfish pain-avoidance and pleasure-seeking are necessary for a DN to devote its limited neuronal and time resource to environmental events that are important according to its selfish pain avoidance and pleasure seeking.

Transparency of decision mechanisms appears to be clearer in terms of the DP of DN compared to rule-based task-specific rules of traditional assistive robots, as the decision mechanisms for muddy tasks are too complex to be handcrafted sufficiently. In this regard, we may consider why market economy is better than any planned economy since economy is also a highly muddy subject.

In summary, if we want our assistive robots to deal with muddy problems as human assistants in human muddy environments, "closed-skull" appears to be a necessary new condition (not a sufficient condition) so that the human programmer is forbidden to use his grossly sub-optimal intervention to compromise the optimality of the emergent representations since such an intervention leads to a higher brittleness. Of course, from this small dialogue we can see that this new necessary condition has not been widely accepted yet. I hope that this dialog is useful for us to continue this debate.

References:

[1] Cole, M. and Cole, S. R. *The Development of Children*, 3rd ed., Freeman, New York, 1996.
[2] Kandel E. R., and Schwartz, J. H. and Jessell, T. M. and Siegelbaum, S. and Hudspeth, A. J. *Principles of Neural Science*, 5th ed., McGraw-Hill, New York, 2012.
[3] Weng, J. Why Have We Passed `Neural Networks Do not Abstract Well?, *Natural Intelligence: the INNS Magazine*, 1(1), pp. 13-22, 2011.
[4] Weng, J. Symbolic Models and Emergent Models: A Review, *IEEE Transactions on Autonomous Mental Development,* 4(1), pp.29-53, 2012.

Dialog Initiation

How are Grammatical Constructions Linked to Embodied Meaning Representations?

Peter Ford Dominey
CNRS & INSERM U846
Stem Cell and Brain Research Institute
Robot Cognition Laboratory
Peter.dominey@insem.fr

There is a strong tendency in modern cognitive neuroscience to adopt the perspective that the comprehension of meaning is achieved in the brain through the mental simulation of that meaning. In this embodied meaning context, brain networks involved in the active perception of an event would also participate in the re-presentation of such an event when reading a sentence describing that event. A wealth of behavioral and neurophysiological studies, including those arguing for a human "mirror system" support this position. These mental simulations have been referred to in a number of contexts including Johnson-Laird's mental models [1-3], Barsalou's perceptual symbol systems [4, 5], and situated simulation models [6]. Major open issues remain, however.

The first issue concerns the development of language: how children learn to use grammar that allows the specification of the temporal unfolding of events in simulation? Does the progressive increase in the complexity of grammatical constructions that are used in development correspond to a developing capability to mentally represent? Do these capacities for language and simulation co-develop, and is there a dependency relation? We have attempted to address these issues of co-development [7, 8], arguing for a form of "conceptual bootstrapping" [9], where the conceptual system provides structure on which the grammatical system is built [10, 11]. But what if man's unusually developed simulation capability owes part of its power to language? It is also potentially problematic that our method of linking language to meaning required a propositional link between the perceptual system and the language system.

This raises a second important question concerning how these simulations are linked to language. In particular, is there a direct link between language and simulation? Or must the target of simulation be coded symbolically? Can the language system directly access simulations via multimodal convergence zones [12]?

A third major open issue concerns the details of how these simulations are managed, the unrolling of these simulations in time, and how grammatical structure orchestrates the "internal film" of mental simulation. Madden et al. [13] hold that "language allows the speaker to "direct the film", to precisely control the initiation, unfolding and termination of appropriate simulations in the mind of the listener, through precise grammatical mechanisms that have evolved for this purpose". Bergen and Chang [14] have begun to address how verb aspect can influence simulation.

Thus, it is likely that comprehension involves a coordinated cooperation between amodal linguistic representations, and modal simulations, as proposed in the Language and Situtated Simulation theory of Barsalou [15]. This raises important questions for future research [16]: How does grammar connect to mental simulations, to turn them on and off, to fast forward them? How does grammatical aspect (which allows the specification of the completion status of actions) interact with situated simulations?

References:

[1] Johnson-Laird, P. Mental Models in Cognitive Science, *Cognitive Science,* 4, pp. 71-115, 1980.
[2] Garnham, A., Oakhill, J. and Johnson-Laird, P. N. Referential continuity and the coherence of discourse, *Cognition,* 11, pp. 29-46, Jan 1982.
[3] Johnson-Laird, P. N. *Mental models: Towards a cognitive science of language, inference, and consciousness,* Harvard University Press, 1983.
[4] Barsalou, L. W. Perceptual symbol systems, *Behav Brain Sci,* 22, pp. 577-609; discussion 610-60, Aug 1999.
[5] Barsalou, L. W., Simmons, W. Barbey, A. and Wilson, C. Grounding conceptual knowledge in modality-specific systems, *Trends Cogn Sci,* 7, pp. 84-91, Feb 2003.
[6] Barsalou, L. W. Simulation, situated conceptualization, and prediction, *Philos Trans R Soc Lond B Biol Sci,* 364, pp. 1281-9, May 12 2009.
[7] Dominey P. and Boucher, J. Learning to talk about events from narrated video in a construction grammar framework, *Artificial Intelligence,* 167, pp. 31-61, 2005.
[8] Dominey P. and Boucher, J. Developmental stages of perception and language acquisition in a perceptually grounded robot, *Cognitive Systems Research,* 6, pp. 243-259, 2005.
[9] Dominey, P. F. Conceptual grounding in simulation studies of language acquisition, *Evolution of Communication,* 4, pp. 57-85, 2002.
[10] Dominey, P. F. Conceptual grounding in simulation studies of language acquisition, *Evolution of Communication,* 4, pp. 57-85, 2000.
[11] Dominey, P. F. A conceptuocentric shift in the characterization of language: Comment on Jackendoff, *Behavioral and Brain Sciences,* 26, pp. 674-674, 2003.
[12] Meyer, K. and Damasio, A. Convergence and divergence in a neural architecture for recognition and memory, *Trends Neurosci,* 32, pp. 376-82, Jul 2009.
[13] Madden, C. Hoen, M. and Dominey, P. F. A cognitive neuroscience perspective on embodied language for human-robot cooperation, *Brain Lang,* 112, pp. 180-8, Mar 2010.
[14] Bergen, B. and Chang, N. Embodied construction grammar in simulation-based language understanding, *Construction grammars: Cognitive grounding and theoretical extensions,* pp. 147-190, 2005.
[15] Barsalou, L. W., Santos, A. Simmons, W. and Wilson, C. Language and simulation in conceptual processing, *Symbols, embodiment, and meaning,* pp. 245-283, 2008.
[16] Pezzulo, G., Barsalou, L. W., Cangelosi, A. Fischer, M. McRae, K. and Spivey, M. J. The mechanics of embodiment: a dialog on embodiment and computational modeling, *Front Psychol,* 2, p. 5, 2011.

Call for Participation

IEEE International Conference on Development and Learning and Epigenetic Robotics (ICDL-EpiRob), 2013

Date: August 18-22, 2013
Location: Osaka City Central Public Hall, Osaka, Japan

Web: http://www.icdl-epirob.org

General chairs:
Yukie Nagai, Osaka University
Angelo Cangelosi, Plymouth University
Chen Yu, Indiana University

Program Chairs:
Jun Tani, KAIST
Giorgio Metta, IIT

Invited Keynote Speakers:
Anne Fernald, Stanford University
Herbert Jaeger, Jacobs University
Yasuo Kuniyoshi, University of Tokyo
Ichiro Tsuda, Hokkaido University

The past decade has seen the emergence of a new scientific field that studies how intelligent biological and artificial systems develop sensorimotor, cognitive and social abilities, over extended periods of time, through dynamic interactions with their physical and social environments. This field lies at the intersection of a number of scientific and engineering disciplines including Neuroscience, Developmental Psychology, Developmental Linguistics, Cognitive Science, Computational Neuroscience, Artificial Intelligence, Machine Learning, and Robotics. Various terms have been associated with this new field such as Autonomous Mental Development, Epigenetic Robotics, Developmental Robotics, etc., and several scientific meetings have been established. The two most prominent conference series of this field, the International Conference on Development and Learning (ICDL) and the International Conference on Epigenetic Robotics (EpiRob), are now joining forces and invite submissions for a joint meeting in 2013, to explore and extend the interdisciplinary boundaries of this field.

Call for Paper

IEEE TRANSACTIONS ON
AUTONOMOUS MENTAL DEVELOPMENT

IEEE TAMD author-pay Open Access option

IEEE TAMD is a hybrid journal allowing either traditional publications or author-pay Open Access (OA) publications. The OA option, if selected, enables unrestricted public access to the article via IEEE Xplore. The OA option will be offered to the author at the time the manuscript is submitted. If selected, the OA fee must be paid before the article is published in the journal. IEEE currently offers the discounted rate of $1,750 per article. The traditional option, if selected, enables access to all qualified subscribers and purchasers via IEEE Xplore. No OA payment is required. The IEEE peer review standard of excellence is applied consistently to all submissions. All accepted articles will be included in the print issue mailed to subscribers.

- Zhengyou Zhang, Editor in Chief of IEEE Transactions on Autonomous Mental Development

IEEE TAMD Table of Contents

Volume 4, Issue 4, December 2012

A Unified Account of Gaze Following

Jasso, H.; Triesch, J.; Deák, G.; Lewis, J.M. Page(s): 257 - 272 (PDF)

Abstract: Gaze following, the ability to redirect one's visual attention to look at what another person is seeing, is foundational for imitation, word learning, and theory-of-mind. Previous theories have suggested that the development of gaze following in human infants is the product of a basic gaze following mechanism, plus the gradual incorporation of several distinct new mechanisms that improve the skill, such as spatial inference, and the ability to use eye direction information as well as head direction. In this paper, we offer an alternative explanation based on a single learning mechanism. From a starting state with no knowledge of the implications of another organism's gaze direction, our model learns to follow gaze by being placed in a simulated environment where an adult caregiver looks around at objects. Our infant model matches the development of gaze following in human infants as measured in key experiments that we replicate and analyze in detail.

A Developmental Approach to Structural Self-Organization in Reservoir Computing

Jun Yin; Yan Meng; Yaochu Jin Page(s): 273 - 289 (PDF)

Abstract: Reservoir computing (RC) is a computational framework for neural network based information processing. Little work, however, has been conducted on adapting the structure of the neural reservoir. In this paper, we propose a developmental approach to structural self-organization in reservoir computing. More specifically, a recurrent spiking neural network is adopted for building up the reservoir, whose synaptic and structural plasticity are regulated by a gene regulatory network (GRN). Meanwhile, the expression dynamics of the GRN is directly influenced by the activity of the neurons in the reservoir. We term this proposed model as GRN-regulated self-organizing RC (GRN-SO-RC). Contrary to a randomly initialized and fixed structure used in most existing RC models, the structure of the reservoir in the GRN-SO-RC model is self-organized to adapt to the specific task using the GRN-based mechanism. To evaluate the proposed model, experiments have been conducted on several benchmark problems widely used in RC models, such as memory capacity and nonlinear auto-regressive moving average. In addition, we apply the GRN-SO-RC model to solving complex real-world problems, including speech recognition and human action recognition. Our experimental results on both the benchmark and real-world problems demonstrate that the GRN-SO-RC model is effective and robust in solving different types of problems.

Context-Based Bayesian Intent Recognition

Kelley, R.; Tavakkoli, A.; King, C.; Ambardekar, A.; Nicolescu, M.; Nicolescu, M. Page(s): 215 - 225 (PDF)
Abstract: One of the foundations of social interaction among humans is the ability to correctly identify interactions and infer the intentions of others. To build robots that reliably function in the human social world, we must develop models that robots can use to mimic the intent recognition skills found in humans. We propose a framework that uses contextual information in the form of object affordances and object state to improve the performance of an underlying intent recognition system. This system represents objects and their affordances using a directed graph that is automatically extracted from a large corpus of natural language text. We validate our approach on a physical robot that classifies intentions in a number of scenarios..

Predicting Visual Stimuli From Self-Induced Actions: An Adaptive Model of a Corollary Discharge Circuit

Ruesch, J.; Ferreira, R.; Bernardino, A. Page(s): 290 - 304 (PDF)
Abstract: Neural circuits that route motor activity to sensory structures play a fundamental role in perception. Their purpose is to aid basic cognitive processes by integrating knowledge about an organism's actions and to predict the perceptual consequences of those actions. This work develops a biologically inspired model of a visual stimulus prediction circuit and proposes a mathematical formulation for a computational implementation. We consider an agent with a visual sensory area consisting of an unknown rigid configuration of light-sensitive receptive fields which move with respect to the environment and according to a given number of degrees of freedom. From the agent's perspective, every movement induces an initially unknown change to the recorded stimulus. In line with evidence collected from studies on ontogenetic development and the plasticity of neural circuits, the proposed model adapts its structure with respect to experienced stimuli collected during the execution of a set of exploratory actions. We discuss the tendency of the proposed model to organize such that the prediction function is built using a particularly sparse feedforward network which requires a minimum amount of wiring and computational operations. We also observe a dualism between the organization of an intermediate layer of the network and the concept of self-similarity.

Human-Recognizable Robotic Gestures

Cabibihan, J. ; Wing-Chee So ; Pramanik, S. Page(s): 305 - 3143 (PDF)
Abstract: For robots to be accommodated in human spaces and in daily human activities, robots should be able to understand messages from their human conversation partner. In the same light, humans must also understand the messages that are being communicated to them by robots, including nonverbal messages. We conducted a Web-based video study wherein participants interpreted the iconic gestures and emblems produced by an anthropomorphic robot. Out of the 15 robotic gestures presented, we found 6 that can be accurately recognized by the human observer. These were nodding, clapping, hugging, expressing anger, walking, and flying. We review these gestures for their meaning from literature on human and animal behavior. We conclude by discussing the possible implications of these gestures for the design of social robots that are able to have engaging interactions with humans.

Intrinsic Motivation and Introspection in Reinforcement Learning

Merrick, K.E. Page(s): 315 - 329 (PDF)
Abstract: Incorporating intrinsic motivation with reinforcement learning can permit agents to independently choose, which skills they will develop, or to change their focus of attention to learn different skills at different times. This implies an autonomous developmental process for skills in which a skill-acquisition goal is first identified, then a skill is learned to solve the goal. The learned skill may then be stored, reused, temporarily ignored or even permanently erased. This paper formalizes the developmental process for skills by proposing a goal-lifecycle using the option framework for motivated reinforcement learning agents. The paper shows how the goal-lifecycle can be used as a basis for designing motivational state-spaces that permit agents to reason introspectively and autonomously about when to learn skills to solve goals, when to activate skills, when to suspend activation of skills or when to delete skills. An algorithm is presented that simultaneously learns: 1) an introspective policy mapping motivational states to decisions that change the agent's motivational state, and 2) multiple option policies mapping sensed states and actions to achieve various domain-specific goals. Two variations of agents using this model are compared to

motivated reinforcement learning agents without introspection for controlling non-player characters in a computer game scenario. Results show that agents using introspection can focus their attention on learning more complex skills than agents without introspection. In addition, they can learn these skills more effectively.

Model-Free Reinforcement Learning of Impedance Control in Stochastic Environments

Stulp, F.; Buchli, J.; Ellmer, A.; Mistry, M.; Theodorou, E.A.; Schaal, S. Page(s): 330 - 341 (PDF)

Abstract: For humans and robots, variable impedance control is an essential component for ensuring robust and safe physical interaction with the environment. Humans learn to adapt their impedance to specific tasks and environments; a capability which we continually develop and improve until we are well into our twenties. In this article, we reproduce functionally interesting aspects of learning impedance control in humans on a simulated robot platform. As demonstrated in numerous force field tasks, humans combine two strategies to adapt their impedance to perturbations, thereby minimizing position error and energy consumption: 1) if perturbations are unpredictable, subjects increase their impedance through co-contraction; and 2) if perturbations are predictable, subjects learn a feed-forward command to offset the perturbation. We show how a 7-DOF simulated robot demonstrates similar behavior with our model-free reinforcement learning algorithm PI2, by applying deterministic and stochastic force fields to the robot's end-effector. We show the qualitative similarity between the robot and human movements. Our results provide a biologically plausible approach to learning appropriate impedances purely from experience, without requiring a model of either body or environment dynamics. Not requiring models also facilitates autonomous development for robots, as pre-specified models cannot be provided for each environment a robot might encounter.

Volume 5, Issue 1, March 2013

The Coordinating Role of Language in Real-Time Multimodal Learning of Cooperative Tasks

Petit, M. ; Lallee, S. ; Boucher, J.-D. ; Pointeau, G. ; Cheminade, P. ; Ognibene, D. ; Chinellato, E. ; Pattacini, U. ; Gori, I. ; Martinez-Hernandez, U. ; Barron-Gonzalez, H. ; Inderbitzin, M. ; Luvizotto, A. ; Vouloutsi, V. ; Demiris, Y. ; Metta, G. ; Dominey, P.F. Page(s): 3 - 17 (PDF)

Abstract: One of the defining characteristics of human cognition is our outstanding capacity to cooperate. A central requirement for cooperation is the ability to establish a “shared plan”-which defines the interlaced actions of the two cooperating agents-in real time, and even to negotiate this shared plan during its execution. In the current research we identify the requirements for cooperation, extending our earlier work in this area. These requirements include the ability to negotiate a shared plan using spoken language, to learn new component actions within that plan, based on visual observation and kinesthetic demonstration, and finally to coordinate all of these functions in real time. We present a cognitive system that implements these requirements, and demonstrate the system's ability to allow a Nao humanoid robot to learn a nontrivial cooperative task in real-time. We further provide a concrete demonstration of how the real-time learning capability can be easily deployed on a different platform, in this case the iCub humanoid. The results are considered in the context of how the development of language in the human infant provides a powerful lever in the development of cooperative plans from lower-level sensorimotor capabilities.

A Survey of the Ontogeny of Tool Use: From Sensorimotor Experience to Planning

Guerin, F. ; Kruger, N. ; Kraft, D. Page(s): 18 - 45 (PDF)

Abstract: In this paper, we review current knowledge on tool use development in infants in order to provide relevant information to cognitive developmental roboticists seeking to design artificial systems that develop tool use abilities. This information covers: 1) sketching developmental pathways leading to tool use competences; 2) the characterization of learning and test situations; 3) the crystallization of seven mechanisms underlying the developmental process; and 4) the formulation of a number of challenges and recommendations for designing artificial systems that exhibit tool use abilities in complex contexts.

Learning Information Acquisition for Multitasking Scenarios in Dynamic Environments

Karaoguz, C. ; Rodemann, T. ; Wrede, B. ; Goerick, C. Page(s): 46 - 61 (PDF)

Abstract: Real world environments are so dynamic and unpredictable that a goal-oriented autonomous system performing a set of tasks repeatedly never experiences the same situation even though the task routines are the same. Hence, manually designed solutions to execute such tasks are likely to fail due to such variations. Developmental approaches seek to solve this problem by implementing local learning mechanisms to the systems that can unfold capabilities to achieve a set of tasks through interactions with the environment. However, gathering all the information available in the environment for local learning mechanisms to process is hardly possible due to limited resources of the system. Thus, an information acquisition mechanism is necessary to find task-relevant information sources and applying a strategy to update the knowledge of the system about these sources efficiently in time.

A modular systems approach may provide a useful structured and formalized basis for that. In such systems different modules may request access to the constrained system resources to acquire information they are tuned for. We propose a reward-based learning framework that achieves an efficient strategy for distributing the constrained system resources among modules to keep relevant environmental information up to date for higher level task learning and executing mechanisms in the system. We apply the proposed framework to a visual attention problem in a system using the iCub humanoid in simulation.

A Spike-Based Model of Neuronal Intrinsic Plasticity

Chunguang Li ; Yuke Li. Page(s): 62 - 73 (PDF)

Abstract: The discovery of neuronal intrinsic plasticity (IP) processes which persistently modify a neuron's excitability necessitates a new concept of the neuronal plasticity mechanism and may profoundly influence our ideas on learning and memory. In this paper, we propose a spike-based IP model/adaptation rule for an integrate-and-fire (IF) neuron to model this biological phenomenon. By utilizing spikes denoted by Dirac delta functions rather than computing instantaneous firing rates for the time-dependent stimulus, this simple adaptation rule adjusts two parameters of an individual IF neuron to modify its excitability. As a result, this adaptation rule helps an IF neuron to keep its firing activity in a relatively “low but not too low” level and makes the spike-count distributions computed with adjusted window sizes similar to the experimental results.

Autonomous and Interactive Improvement of Binocular Visual Depth Estimation through Sensorimotor Interaction

Mann, T.A. ; Yunjung Park ; Sungmoon Jeong ; Minho Lee ; Yoonsuck Choe. Page(s): 74 - 84 (PDF)

Abstract: We investigate how a humanoid robot with a randomly initialized binocular vision system can learn to improve judgments about egocentric distances using limited action and interaction that might be available to human infants. First, we show how distance estimation can be improved autonomously. We consider our approach to be autonomous because the robot learns to accurately estimate distance without a human teacher providing the distances to training targets. We find that actions that, in principle, do not alter the robot's distance to the target are a powerful tool for exposing estimation errors. These errors can be used to train a distance estimator. Furthermore, the simple action used (i.e., neck rotation) does not require high level cognitive processing or fine motor skill. Next, we investigate how interaction with humans can further improve visual distance estimates. We find that human interaction can improve distance estimates for far targets outside of the robot's peripersonal space. This is accomplished by extending our autonomous approach above to integrate additional information provided by a human. Together these experiments suggest that both action and interaction are important tools for improving perceptual estimates.

Editor : Pierre-Yves Oudeyer, Inria Editorial Assistants : Adrien Baranes and Paul Spears

The AMD Newsletter is financially supported by the IEEE CIS and Inria