Editorial

Language acquisition has probably been the centre of the most profound, and yet unsolved, scientific debates in cognitive sciences in the 20[th] century. Computational and robotic models elaborated in the recent years aim at overcoming a number of conceptual barriers in this debate, because they help to naturalize these questions and allow us for the first time to confront theories to reality. This issue of the newsletter features a dialog, initiated by Angelo Cangelosi, over the symbol grounding problem, in relation to advances in robotics, machine learning and artificial intelligence. Central scientific actors of the history of this problem have responded. Stevan Harnad, Luc Steels, Aaron Sloman, Stephen Cowley, Vincent Müller, Carol Madden, Peter Ford Dominey and Stéphane Lallée present their own, sometimes very constrastive, views, and we see that in spite of some clear progress, a lot of work is in front of us, both conceptually and in terms of robotic experiments.

Then, a novel call for dialog is proposed by Gianluca Baldassarre and Marco Mirolli and relates to the challenges of cumulative learning : « What are the key open challenges for understanding autonomous cumulative learning of skills? ». Interested researchers are welcome to submit a response (contact pierre-yves.oudeyer@inria.fr) by September 1[st], 2010. The length of each response must be between 300 and 500 words (including references).

I take the opportunity of this editorial to congratulate Zhengyou Zhang for his work as a chair of the AMD TC, which has been essential for the growth of our community, and in particular his strong support for this newsletter. I also welcome Minoru Asada as the new AMD TC chair, and I am sure his vision will help all of us to increase even more the momentum of computational developmental sciences.

- Pierre-Yves Oudeyer, INRIA, Editor

Message from the Past Chair of IEEE AMD Technical Committee

It is my great pleasure to welcome Prof. Minoru Asada, Osaka University, Japan, as the new chair of the AMD TC. Minoru and I knew each other more than 20 years ago when we were both doing computer vision. He recognized very early the importance of developmental mechanisms, and started working on developmental robotics early 90's. He is also one of the main organizers of the RoboCup competitions. In the last few years, he has been helping me organize the AMD TC activities as a Vice Chair. I trust that our TC will further grow under this new leadership.

I have been privileged to enjoy the support and encouragement of the AMD community, which has led to the continuous growth of our TC and to the establishment of the new IEEE Transactions on Autonomous Mental Development (TAMD).
In particular, besides Minoru, I would like to thank Juyang (John) Weng (AMD TC Past Chair), James (Jay) L. McClelland (Former President of the Cognitive Science Society), Pierre-Yves Oudeyer (Editor-in-Chief of the AMD TC Newsletter), and the IEEE CIS leadership team for their vital roles in growing our community.

IEEE TAMD is still in its infancy. It needs substantial care to develop healthily. Please submit your own papers, help review submissions, and advertise TAMD to the widest possible audience. Your involvement in TAMD activities is crucial.
Thank you!

- Zhengyou Zhang, Past Chair of the AMD TC (2007-2009)

Message from the New Chair of IEEE AMD Technical Committee

First of all, as a new chair of the AMD TC, I would like to thank the former chair, Dr. Zhengyou Zhang for his role in leading our community and his initiation of a new IEEE transaction on AMD. Also, on the behalf of the AMD TC, I appreciate your continuous supports for our activities. The AMD TC is the largest one among IEEE CIS TCs, (see http://ieee-cis.org/technical/welcome/). Therefore, our activities are also expected to be among the most exciting. In order to cope with the expectations, I would like to enrich and extend our activities, especially ICDL with more people from different disciplines, and co-location of ICDL with other related and major conferences such as EpiRob and IEEE CIS conferences. Furthermore, we should have a systematic way to encourage the authors of ICDL to submit their extended papers to TAMD. We need your continuous supports for these extensions. Thank you for your cooperation in advance.

-Minoru Asada, New Chair of the AMD TC

Dialog Column

The Symbol Grounding Problem Has Been Solved: Or Maybe Not?

Angelo Cangelosi
Adaptive Behaviour & Cognition Lab
Centre for Robotics and Neural Systems, University of Plymouth, UK.

The issue of symbol grounding is of crucial importance to the community of developmental robotics as in the last decades there has been a tremendous increase in new models of language learning, and evolution, in cognitive agents and robots. Although in the literature on AI, cognitive science and philosophy there has been extensive discussion about the symbol grounding problem, there are still quite different views on its importance, ranging from "symbolic" approaches that practically ignore the cognitive significance of such an issue (e.g. Fodor 1983), to "embodied" approaches that acknowledge its importance, but suggest that the problem has practically been solved (Steels 2008).

To assess better the current state of the art on the Symbol Ground Problem, and identify the research challenges and issues still pending, I will use the definition and discussion of the problem originally given by Stevan Harnad in the seminal 1990 article "The Symbol Grounding Problem". Harnad explains that the symbol grounding problem refers to the capability of natural and artificial cognitive agents to acquire an intrinsic link (autonomous, we would say in nowadays robotics terminology) between internal symbolic representations and some referents in the external word or internal states. In addition, Harnad explicitly proposes a definition of a symbol that requires the existence of logical links (e.g. syntactic) between the symbols themselves. It is thanks to these inter-symbol links, its associated symbol manipulation processes, and the symbol grounding transfer mechanism (Cangelosi & Riga 2006) that a symbolic system like human language can exist. The symbol-symbol link is the main property that differentiates a real symbol from an index, as in Peirce's semiotics. These symbolic links also support the phenomena of productivity and generativity in language and contribute to the grounding of abstract concepts and symbols (Barsalou 1999). Finally, an important component of the symbol grounding problem is the social and cultural dimension, that is the role of social interaction in the sharing of symbols (a.k.a. the external/social symbol grounding problem, as in Cangelosi 2006; Vogt 1997).

To summarize, we can say that there are three sub-problems in the development of a grounded symbol system:
1. how can a cognitive agent autonomously link symbols to referents in the world such as objects, events and internal and external states?
2. how can an agent autonomously create a set of symbol-symbol relationships and the associated transition from an indexical system to a proper symbol system?
3. how can a society of agents autonomously develop a shared set of symbols?

I agree with Steels (2008) that much has been done on the robotics and cognitive modeling of the symbol grounding problem when we consider the two sub-problems (1) and (3): "we now understand enough to create experiments in which groups of agents self-organize symbolic systems that are grounded in their interactions with the world and others" (Steels 2008: page 240). But, as Steels also acknowledges, it is also true that we do not yet have a full understanding of all mechanisms in grounding, such as on the nature, role and making of internal symbolic representations.

Dialog Column

As for the sub-problem (2), i.e. the transition from a communication systems based on indices (e.g. labels, animal communication, early child language learning) to that of a full symbolic system (e.g. adult human languages), I believe that the problem has not really been solved at all, and much needs to be done. Most computational models of syntactic learning and evolution use a symbolic approach to this problem, i.e. by assuming the pre-existence of semantic and syntactic categories in the agent's cognitive system. This is however in contrast with the grounding principles.

I invite my colleagues to comment on the state of the art on the symbol grounding problem in developmental robotics models of communication and language, and on their view on the importance (or not!) of the symbol grounding problem. I suggest below some open challenges for future research that I believe are crucial for our understanding of the symbol grounding phenomena, and I welcome suggestions for other important, unsolved challenges in this field:

1. Is the symbol grounding problem, and the three sub-problems as identified above, still a real crucial issue in cognitive robotics research? And if the problem appears to have been solved, as some have suggested, why is it that so far we have failed at building robots that can learn language like children do?

2. What are the developmental and evolutionary processes that lead to the transition from an indexical communication system to a full symbolic system such as language? Is there a continuum between indices (labels) and symbols (words), or is the transition qualitative and sudden? What known phenomena in language origins theories, and in developmental studies, should be included in developmental and evolutionary robotics model of language?

3. Notwithstanding the importance of the grounding problem, there are still various approaches in the agent/robot language learning/evolution literature that practically ignore the process of grounding and use a symbolic-only approach to the definition of meanings and words. Do these symbolic approaches really give an important contribution to our understanding of human cognition, or should all models of language learning be based solely on grounding mechanisms?

4. Does cognitive development really play an important role in symbol grounding and acquisition, or is it just an epiphenomenon of no crucial importance to the understanding of human cognition? Some key findings and experiments show that infants have strong specific biases that allow them to learn language very easily. And most attempts at building robots without these biases have failed so far to learn realistically complex concepts/semantic categories. Is the symbol grounding problem just a matter of using and identifying such biases in robotics language models?

5. What kind of robotic experiment would constitute a real breakthrough to advance the debate on symbol grounding, and what kind of principle and ideas are still unexplored?

6. What are the properties and differences of internal representations beyond both indexical and symbolic systems? Or are representation issues not really crucial, as a pure sensorimotor modelling approach would not require any internal representation capability?

7. How can we model the grounding of abstract concepts such as beauty, happiness, time. Or is the grounding approach inconsistent with the study of higher-order symbolic capabilities?

8. What are the grounding components in the acquisition and use of function words (such as verb preposition "to", as in verbs "to go", "if", "the"), of number concepts/words, and of morphology and other syntactic properties.

9. How can we model the grounding phenomena studies through empirical investigations of language embodiment (Barsalou 1999; Glenberg & Kaschak 2002; Pecher & Zwaan 2005)?

References:

- Barsalou, L. (1999) Perceptual symbol systems. *Behavioral and Brain Sciences*, 22: 577-609.

- Cangelosi A. (2006) The grounding and sharing of symbols. *Pragmatics and Cognition*, 14(2), 275-285

- Cangelosi A, Riga T (2006) An embodied model for sensorimotor grounding and grounding transfer: Experiments with epigenetic robots, *Cognitive Science*, 30(4), 673-689

- Fodor, J.A. (1983) *The modularity of mind: An essay on faculty psychology*. Cambridge, MA: Bradford Books, MIT Press.

- Glenberg A., Kaschak M. (2002). Grounding language in action. *Phychonomic Bulletin & Review*, 9(3):558–565

- Harnad, S. (1990) The Symbol Grounding Problem. *Physica* D 42: 335-346.

- Pecher D., Zwaan R.A. (2005). *Grounding Cognition: The Role of Perception and Action in Memory, Language, and Thinking*: Cambridge University Press.

- Steels, L. (2008) The symbol grounding problem has been solved. So what's next? In de Vega, M., (eds.), *Symbols and Embodiment: Debates on Meaning and Cognition*. Oxford: Oxford University Press. pp. 223-244.

- Vogt P. (2002). Physical symbol grounding. *Cognitive Systems Research*, 3(3):429--457.

Dialog Column

Symbol Grounding Problem: Turing-Scale Solution Needed

Stevan Harnad
Chaire de recherche du Canada en sciences cognitives, Université du Québec à Montréal, Canada,
School of Electronics and Computer Science, University of Southampton, UK

Toys. The symbol grounding problem is the problem of causally connecting symbols inside an autonomous system to their referents in the external world without the mediation of an external interpreter. The only way to avoid triviality, however, is to ensure that the symbols in question, their referents in the world, and the dynamic capacities of the autonomous system interacting with the world are nontrivial. Otherwise a toy robot, with exactly two symbols – go/stop – is "grounded" in a world where it goes until it bumps into something and stops.

Turing. From the very outset, the symbol grounding problem – which was inspired and motivated by Searle's Chinese Room Argument – was based on the Turing Test, and hence on a system with *full, human-scale linguistic capacity*. So it is the words of a full-blown natural language (not all of them, but the ones that cannot be grounded by definition in the others) that need to be connected to their referents in the world. Have we solved that problem? Certainly not. Nor do we have a robot with Turing-scale capacities, either symbolic or sensorimotor (with the former grounded – embodied -- in the latter). Designing or reverse-engineering an autonomous system with this Turing-scale robotic and linguistic capacity -- and thereby causally explaining it -- is the ultimate goal of cognitive science. (Grounding, however, is not the same as meaning; for that we would also have to give a causal explanation of consciousness, i.e., feeling, and that, unlike passing the Turing Test, is not just hard but hopeless.)

Totality. Grounded robots with the sensorimotor and learning capacities of subhuman animals might serve as waystations, but the gist of the Turing methodology is to avoid being fooled by arbitrary fragments of performance capacity. Human language provides a natural totality. (There are no partial languages, in which you can say this, but not that.) We are also extremely good at "mind-reading" human sensorimotor performance capacity for tell-tale signs of mindlessness; it is not clear how good we are with animals (apart perhaps from the movements and facial expressions of the higher mammals).

Terms. There are certain terms (or concepts) I have not found especially useful. It seems to me that real objects -- plus (internal or external) (1) analogs or iconic copies of objects (similar in shape) and (2) arbitrary-shaped symbols in a formal symbol system (such as "x" and "=", or the words in a language, apart from their iconic properties), systematically interpretable as referring to objects -- are entities enough. Peirce's "icon/index/symbol" triad seems one too many. Perhaps an index is just a symbol in a toy symbol system. In a formal symbol system the links between symbols are syntactic whereas the links between internal symbols and the external objects that they are about are sensorimotor (hence somewhat iconic). And inasmuch as symbols inside a Turing-scale robot are linked to object categories rather than to unique (one-time, one-place) individuals, all categories are abstract (being based on the extraction of sensorimotor invariants), including, of course, the category "symbol." The rest is just a matter of degree-of-abstraction. Even icons are abstract, inasmuch as they are neither identical nor co-extensive with the objects they resemble. There are also two sorts of productivity or generativity: syntactic and semantic. The former is just formal; the latter is natural language's power to express any and every truth-valued proposition.

Talk. Yes, language is fundamentally social in that it would never have bothered to evolve if we had been solitary monads (even monads born mature: no development, just cumulative learning capacity). But the nonsocial environment gives enough corrective feedback for us to learn categories. Agreeing on what to call them is trivial. What is not trivial is treating symbol strings as truth-valued propositions that describe or define categories.

References:

- Harnad, S. (2003) Can a Machine Be Conscious? How? *Journal of Consciousness Studies* 10(4-5): 69-75. http://eprints.ecs.soton.ac.uk/7718/
- Harnad, S. (2005) *To Cognize is to Categorize: Cognition is Categorization*, in Lefebvre, C. and Cohen, H., Eds. *Handbook of Categorization*. Elsevier. http://eprints.ecs.soton.ac.uk/11725/
- Harnad, S. and Scherzer, P. (2007) First, Scale Up to the Robotic Turing Test, Then Worry About Feeling. In *Proceedings of Proceedings of 2007 Fall Symposium on AI and Consciousness*. Washington DC. http://eprints.ecs.soton.ac.uk/14430/

Dialog Column

- Harnad, S. (2008) The Annotation Game: On Turing (1950) on Computing, Machinery and Intelligence. In: Epstein, Robert & Peters, Grace (Eds.) *Parsing the Turing Test: Philosophical and Methodological Issues in the Quest for the Thinking Computer*. Springer
  http://eprints.ecs.soton.ac.uk/7741/
- Picard, O., Blondin-Masse, A., Harnad, S., Marcotte, O., Chicoisne, G. and Gargouri, Y. (2009) Hierarchies in Dictionary Definition Space. In: 23rd Annual Conference on Neural Information Processing Systems: Workshop on Analyzing Networks and Learning With Graphs, 11-12 December 2009, Vancouver BC (Canada). http://eprints.ecs.soton.ac.uk/18267/

Work on Symbol Grounding Now Needs Concrete Experimentation

Luc Steels
Vrije Universiteit Brussel (VUB), Brussels, Belgium
and Sony Computer Science Laboratory, Paris, France.

A few years ago, I wrote a paper called "The symbol grounding problem has been solved. So What's Next?" (Steels, 2008) based on a talk I gave in 2005 at the very interesting Garachico workshop on "Symbols, Embodiment, and Meaning" organized by Manuel de Vega, Arthur Glenberg and Arthur Graesser. Angelo Cangelosi, who himself has made several contributions in this area (Cangelosi and Harnard, 2000), responded to this paper calling for further discussion within the developmental robotics community and I of course welcome such dialogs which are all too rare in our field. Cangelosi decomposed the symbol grounding problem in three subproblems: (1) How can a cognitive agent autonomously link symbols to referents in the world such as objects, events and internal and external states? (2) How can an agent autonomously create a set of symbol-symbol relationships and the associated transition from an indexical system to a proper symbol system? and (3) How can a society of agents autonomously develop a shared set of symbols? I will take the same decomposition to argue again why I believe the symbol grounding problem has been solved.

Let's start with (2). AI has developed over the past decade a large array of mechanisms for dealing with symbol-symbol relationships. It suffices to open up any textbook, particularly in machine learning, and you will find effective algorithms that will induce symbolic relationships and inferential structures from sets of symbols (usually texts but they can also be more formalized representations of knowledge as in the semantic web). The success of this body of techniques is undeniable when you look at search engines such as Google which are entirely based on them. This problem can therefore be considered solved, or at least sufficiently solved that applications can be built on a large scale.

The value of the philosophical discussion initiated by Searle and Harnad was to point out that there is more to symbol usage than symbol-symbol linking, in particular there is problem (1): How can symbols be grounded in sensory-motor experiences. Searle and Harnad are right, but also this question was already addressed by early AI systems such as Shakey (built in the early seventies by Nilsson, et.al. at SRI (Nilsson, 1984)). The Shakey experiment was unfortunately buried in technical papers which were not accessible to non-experts and its significance has therefore been missed by philosophers. Shakey had a cognitive architecture designed to trigger embodied grounded robot behaviors through natural language commands. Given the state of the art in computing, robotics, visual perception, etc., the response of the robot was slow and the complexity of what could be handled limited, but the fundamental principle how to ground symbols through feature extraction, segmentation, pattern recognition, etc. was clearly demonstrated and many subsequent experiments have confirmed it.

Philosophers undoubtly object that it is the designers that take care of the grounding. Although the robot autonomously carries out grounding, it does not autonomously acquire the competence to do so. This criticism is right, but I believe this problem has been solved as well. What we need to do is to set up cooperative interactions between agents in which symbols are useful. For example, I am in a store and want to get from you (the salesperson) a brown T-shirt (and not the white, green, yellow, or blue one). And suppose that I cannot point to the T-shirt because they are not on display. Then if I have a way to categorize colors and associate names with these categories, and you have a way to see what category I mean and use that to find the T-shirt back in the world, then I can get the T-shirt I want and you are happy because you can sell it to me.

Dialog Column

Similar situations can be set up with robots, for example because they need to draw attention to an object in the world, or get the other one to do an action, or inform each other about an event that took place, etc. The key point is of course that in these experiments neither the categorizations of reality nor the symbols themselves should be supplied by designers. They should be invented by the agents themselves. And this is precisely what the language games experiments, from the first one reported in Steels (1995), have shown. The key idea is to establish a coupling between success in a task (where communication is critical or at least useful) and the grounding and linguistic processes that are potentially able to handle symbols.

Interestingly enough, the solution to (2) implies a solution to (3). Producers and interpreters must negotiate a tacit agreement how they are going to conceptualize the world for language and how they are going to name and express concepts otherwise the interaction will not be successful. The solution, again demonstrated for the first time in Steels (1995) but since then applied on a grand scale in dozens of robotic experiments, is to monitor success in communication and adapt or shift concepts and symbolic conventions based on that so that conceptual and symbolic inventories gradually align. When this is done systematically by all agents for every language game they play, a process of self-organization towards a shared system is set in motion, without the need for the intervention of a human designer to coordinate symbol use. And this solves subproblem (3).

What should we do now? Work on the specifics, particularly if you are interested in mental development: for example, how can concepts involved in the representation of time arise and get symbolized as tense and aspect? How can a system for categorizing the roles of participants in events (agent, patient, etc.) emerge? Why and how can modality systems to express the attitude of speakers towards information arise? And so on. Much work is left to do but chasing a philosophical chimera is not one of them.

References:

- Cangelosi A. & Harnad S. (2000). The adaptive advantage of symbolic theft over sensorimotor toil: Grounding language in perceptual categories. Evolution of Communication 4(1), 117-142.
- Nilsson, Nils J. (1984) Shakey The Robot, Technical Note 323. AI Center, SRI International, 333 Ravenswood Ave., Menlo Park, CA 94025, Apr 1984. (available from http://www.ai.sri.com/shakey/)
- Steels, L. (1995) A self-organizing spatial vocabulary. Artificial Life Journal, 2(3): 319-332.
- Steels, L. (2008) The Symbol Grounding Problem has been solved. So What's Next? In: Glenberg, A., A. Graesser, and M. de Vega (eds) (2008) Symbols, Embodiment and Meaning. Oxford University Press, Oxford. pp. 506-557.

Becoming Symbol Sharers

Stephen Cowley
School of Psychology, University of Hertfordshire,
College Lane, Hatfield, UK

Imagine outsider-agents who, while cared for, are marginal in a population that relies on evolved public symbols. The insider-agents have solved the first and third parts of Cangelosi's puzzle. Not only does the population use shared symbols but, as functioning bodies, their controlling systems enable them to identify public referents. How, then, do outsider-agents solve the problem of becoming symbol sharers?

Inner symbols use contingencies that lack any causal link with the population's external symbols. Outsiders become insiders by drawing on developmental history. By hypothesis, this depends on *coaction*. In Wegner and Sparrow's (2007) terms, each party is influenced by, or acts within, the context of the other's actions: dyads come up with ways of moving – and routines – that cannot be enacted alone. If extended to language, we need two theoretical postulates:

- Total language supports concurrent dynamical and symbolic descriptions.
- In external symbol grounding, outsider-agents use symbolic patterns to gain control over *their own* doings.

Dialog Column

Since dynamics co-occur with symbols, external patterns become *cues to meaning* (Elman, 2009). No mental lexicon can connect vocalization, movement and settings: language spreads across space, time, bodies and brains.

How do infants become symbol sharers? As rewards and contingencies become familiar, norms evoke symbolic patterns. Babies play, feel and learn: by perceiving, babbling and acting predictably, they gain control over their own actions. *Human symbol grounding* integrates routines with both heard and unheard verbal patterns. By 3-4 months babies pick up on cultural norms; by 6-8, they share formats (e.g. bath-time; nappy changing) and by 9-12, they act as if using intentions. While learning is reward-based, public construals sensitise infants to situations. Far from representing utterance-types ('linguistic forms'), dyads use co-regulation. The results shape how children perceive. That is nature's trick: we use external patterns and what we *hear* and *see* to redeploy our bodies and brains (Anderson, 2008). Once we perceive 'words' babies answer – and ask – new questions ("Do you understand?" "What's that?" "Where's the car?"). A shared *view* of language transforms perception, cultural skills and first-person phenomenology. Since words become 'real', there is no need for qualitative neural change. Children can rely on perceiving external symbolic patterns –and talking sensibly about them. By integrating the results with action, children develop new control over their own doings: they become symbol sharers.

Symbol sharing arises from how babies use public construals. Like robots, they elicit real-time understanding: to adopt local practices, they need not know what they are doing. Seen thus, Cangelosi's second question opens up the *problem of coaction*. Can we link inner representations with external symbolic patterns by designing systems that learn as, in real-time, actions are co-regulated? As with babies, dynamics and symbols –speaking and moving –would be resources used to assess and manage *situations*.

References:

- Anderson, M. L. (2008). Circuit sharing and the implementation of intelligent systems. Connection Science, 20/4: 239 – 251.
- Elman, J.L. (2009). On the meaning of words and dinosaur bones: Lexical knowledge without a lexicon. *Cognitive Science*, 33/4: 547–582.
- Wenger, D. & Sparrow, B. (2007). The Puzzle of Coaction. In: *Distributed Cognition and the Will*, D. Ross, D. Spurrett, H. Kinkaid & L.G. Stephens (Eds.) , pp.17-41. MIT Press: Cambridge MA.

Symbol Grounding is Not a Serious Problem. Theory Tethering Is
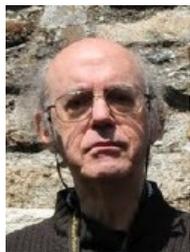
Aaron Sloman
University of Birmingham, School of Computer Science
Birmingham, UK

Misquoting Santayana: "Those who are ignorant of philosophy are doomed to reinvent it – badly".

The "symbol grounding" problem is a reincarnation of the philosophical thesis of concept empiricism: "every concept must either have been abstracted from experience of instances, or explicitly defined in terms of such concepts". Kant refuted this in 1781, roughly by arguing that experience without prior concepts (e.g. of space, time, ordering, causation) is impossible. 20th century philosophers of science added nails to the coffin, because deep concepts of explanatory sciences (e.g. "electron", "gene") cannot be grounded in the claimed way. Symbol grounding is impossible.

Let's start again! All organisms and many machines use information to control behaviour. In simple cases they merely relate motor signals, sensor signals, their derivatives, their statistics, etc. with, or without, adaptation. For a subset of animals and machines that's insufficient, because they need information about things that are not, or even cannot be, sensed, e.g. things and events in the remote past or too far away in space, contents of other minds, future possible events and processes, and situations that can ensue.

Coping with variety and novelty in the environment in a fruitful way requires use of recombinable "elements of meaning" -- concepts, whose vehicles may or may not be usefully thought of as symbols. So a relatively small number of concepts can be recombined to form different percepts, beliefs, goals, plans, questions, conjectures, explanatory theories, etc.. These need not all be encoded as sentences, logical formulae, or other discrete structures.

Dialog Column

E.g. maps, trees, graphs, topological diagrams, pictures of various sorts, 2-D and 3-D models (e.g. of the double helix) are all usable for encoding information structures, for purposes of predicting, recording events, explaining observations, forming questions, controlling actions, communicating, etc. Sometimes new forms of representation need to be invented -- as happens in science, engineering, and, I suspect, in biological evolution and animal development.

How can conceptual units refer? Through their role in useful, powerful theories. In simple cases, a theory is a set of "axioms" with undefined symbols. There will be some portions of reality that are and some that are not models of the axioms. In each model there will be referents for all the undefined symbols. To reduce ambiguity by eliminating some of the models, scientists add "bridging rules" to their theories, linking the theories to experiments and measurements, which support predictions. This partially "tethers" the theory to some portion of reality, though never with final precision, so theories continue being refined internally and re-tethered via new instruments and experimental methods. An infant or toddler learning about its environment, which contains different kinds of stuff, things, events, processes, minds, etc. needs to do something similar, though nobody knows how, and robots are nowhere near this.

Progress requires much clearer philosophical understanding of the requirements and then new self-extending testable designs for meeting the requirements. Most of what has been written about symbol grounding can then be discarded.

More details are available online.
http://www.cs.bham.ac.uk/research/projects/cogaff/talks/#models
http://www.cs.bham.ac.uk/research/projects/cogaff/talks/#babystuff
http://www.cs.bham.ac.uk/research/projects/cosy/papers/#tr0804

The Hard and Easy Grounding Problems

Vincent C. Müller
Department of Humanities & Social Sciences
Anatolia College/ACT

'Grounding' is somehow desirable for robotics, but it is not terribly clear why, how, and for what. Historically, Searle had attacked the claims that a computer can literally be said to understand, or to explain the human ability to understand. His argument was that a computer only does syntactical 'symbol' manipulation and thus a system can not "think, understand and so on solely in virtue of being a computer with the right sort of program" (Searle, 1980, p. 368). Harnad turned this in to the challenge that the symbols in a machine with intentional properties (like understanding or meaning) should be 'grounded' via causal connection to the world (Harnad, 1990). Both of these authors were working against the then default thesis that human cognition is computation. If this is assumed, it follows the same computational functions can be carried out on different hardware, e.g. silicon chips instead of neurons.

On this background, I see four symbol grounding problems:
1. How can a purely computational mind acquire meaningful symbols?
2. How can we get a computational robot to show the right linguistic behaviour?

These two are misleading: The first one cannot be solved (Müller, 2009) and the second one might not involve grounding since it just requires the right output. We need a problem that does not assume computationalism, but goes beyond 'output' towards grounding of public natural language. In analogy to the distinction between a 'hard' and an 'easy' problem about consciousness (Chalmers, 1995), I would suggest:
3. How can we explain and re-produce the behavioural ability and function of meaning in artificial computational agents?
4. How does physics give rise to meaning?

We do not even know how to start on the 'hard problem' in (4): from physics to intentional states and phenomenal consciousness; so we should tackle the 'easy problem' in (3): behaviour and function.

Dialog Column

Cangelosi's three 'sub-problems' are part of this, only that I doubt whether our problem is that of a lonely and fully intelligent agent trying to acquire basic symbols, then to transfer the grounding, and finally to negotiate in communication with other agents. I tend to think that an account of function of natural language and language acquisition in humans will have to account for the function of language in intelligence and for the social function of language.

References:

- Chalmers, D. J. (1995). Facing up to the problem of consciousness. Journal of Consciousness Studies, 2 (3), 200-219.
- Harnad, S. (1990). The symbol grounding problem. Physica D, 42, 335-346.
- Müller, V. C. (2009). Symbol grounding in computational systems: A paradox of intentions. Minds and Machines, 19 (4), 529-541.
- Searle, J. R. (1980). Minds, brains and programs. Behavioral and Brain Sciences, 3, 417-457.

Reply to Dialog: "The Symbol Grounding Problem Has Been Solved: Or Maybe Not?"

Carol J. Madden, Peter Ford Dominey, Stéphane Lallée
CNRS and Inserm
INSERM, Stem Cell & Brain Research Institute, Bron, France

One recent perspective that has been successful in accounting for event representation in humans is that of physicalist models (e.g., Wolff, 2007), in which causes and effects are understood in terms of physical dynamics in the world, such as momentum and impact forces. Furthermore, nonphysical causation (e.g., forcing someone to decide) is understood by analogy to these physical force primitives. This approach can be borrowed for any artificial system having a body and the ability to perceive kinematics and dynamic forces. Of course in the robot, these embodied perceptual and physical primitives will surely be transduced into symbols (steels' c-representations). While this is not necessarily problematic, it should be noted that in human cognition, we can easily follow the link back to grounded perception and action; simulating or imagining what it is like to see a red fire truck or pick up a hot potato, with rich perceptual and motor information (and corresponding activation in perceptual-motor areas). In robot models, once an input is transduced to a symbol, there often remains no real link to the continuous/embodied perceptual-motor inputs. Introducing this type of simulation capacity could provide groundbreaking benefits to meaning representation and event understanding (see Madden et al., in press).

To avoid being trapped at the indexical level, robots need to not only perceive objects and force primitives (e.g., motion, contact), but also integrate these representations to comprehend events. This requires that representational symbols participate in a kind of grammatical or rule-based processing, just as their labels (words) participate in a linguistic grammar, and their referents in the world participate in role-dependent interactions. However, rule-based symbol processing necessitates some sort of categorical variation in the symbols themselves. Therefore, the robot will need different kinds of symbols that correspond to the different types of meaningful things that can occur in the world, such as entities, states, actions, and relations. These different categories of symbols can then be processed in a representational system according to relational rules. Ideally, the differences between symbol types could even be grounded in the different perceptual processes from which they arise.

In our research, the robot can observe, learn, and perform a series of actions during goal-directed interaction with humans. While the objects and actions are recognized through perceptual primitives, it is clear that this yields no real understanding of the events. Through this new approach, instead of representing sequences of actions to build goal directed events, the robot will represent both actions (A) and states (S), so that events are built from rule-based SAS chains of enabling-states, actions, and resulting-states (Lallée et al., submitted). Recoding the perceptual input into these two different kinds of symbols (states, actions) is a first step that opens the door to representing causality and reasoning about goals. For example, a resulting state can be achieved if the enabling state is true and the action is performed, and this resulting state may serve as the enabling state for another stored SAS event.

Dialog Column

Furthermore, this type of representation could be a precursor to understanding intentionality. If a robot is allowed to interact with its environment at random and observe the changes, the robot might observe that a given motor action always produces a given change in state, such as an object being "covered". Eventually, the robot will learn the causal connection between its own action and that resulting state. It can remember this when a task later requires an object to be covered, and initiate the remembered motor action itself to cover the object. By initiating the causal sequence itself, the robot demonstrates a primitive representation of intentionality.

References:

- Lallée, S., Madden, C.J., Hoen, M., & Dominey, P.F. (submitted). Linking Language with Embodied and Teleological Representations of Action for Humanoid Cognition. Frontiers in Neurorobotics.
- Madden, C.J., Hoen, M., & Dominey, P.F. (in press). A cognitive neuroscience perspective on embodied language for human-robot cooperation. Brain and Language.
- Wolff, P. (2007). Representing causation. Journal of Experimental Psychology: General, 136, 82-111

Reply to Commentaries

Angelo Cangelosi
Adaptive Behaviour & Cognition Lab
Centre for Robotics and Neural Systems, University of Plymouth, UK.

I welcome the very useful comments and insights from my esteemed colleagues, all of whom have significantly contributed to our current understanding (and solutions!) of the symbol grounding problem.

I would first like to start with the commentary of my symbol grounding "maestro", Stevan **Harnad**. I welcome his invitation to go beyond using only toy-level solutions to the symbol grounding problem and to aim at Turing-scale experimentation. I believe that one of the most fruitful potential contributions of current developmental robotics advances is to be able to design experiments where a community of robotic agents is able to socially acquire grounded symbolic systems. This is in line with **Steels**'s invitation to address the symbol grounding issues with practical (robotic) experimentation, and with **Sloman**'s encouragement to develop new self-extending testable designs for meeting the requirements developed in philosophical and scientific analyses. **Muller** highlights the distinction between a "hard" version of the problem (i.e. focus on the agent's own intentionality in creating and using symbols) and the "easy" one that should aim at the explanation and re-production of the *behavioral ability* and *function* of meaning in artificial computational agents and robots. These can be investigated through practical experimentations on developmental robotics.

As to the commentaries on the current progress and solutions to the three sub-problems identified in my original target column, the issue that appears to be pending remains that of the sub-problem (2) on the transition from indexical representation to syntactic symbol-symbol representations. **Steels** is right to say that impressive progress has happened on the development of (ungrounded!) symbol-symbol manipulation methods. However, the core problem remains that all these solutions are based on symbolic approaches that use symbols and syntactic relations that are hand-crafted by the researcher, and are therefore not grounded in the agent's own sensorimotor system. Some useful insights are suggested to address such a sub-problem. **Madden**, Lallée and Dominey propose that the study of embodied simulation capacity in robots could provide groundbreaking benefits to meaning representation and event understanding, thus freeing the robot from "being trapped at the indexical level" through the integration of these representations into more articulated, syntactic-like representation systems. **Cowley** also suggests a solution through the focus on the "*problem of coactions*" in child development. In babies, dynamics and symbols –how we speak and move – gradually become resources that could be used in assessing and managing *situations* (which is similar to **Madden'**s et al. mental simulation capability).

Developmental robotics has a great challenge, and opportunity, to allow us to demonstrate through Turing-like robotics experimentations some of the great mysteries of human development, that is how babies are able to acquire autonomously, through interaction with the physical world, own body and their social environment, this symbolic system that we call language.

Dialog Initiation

What Are the Key Open Challenges for Understanding the Autonomous Cumulative Learning of Skills?

Gianluca Baldassarre, Marco Mirolli
Laboratory of Computational Embodied Neuroscience, Istituto di Scienze e Tecnologie
della Cognizione, Consiglio Nazionale delle Ricerche, Rome, Italy

The capacity to autonomously learn a number of different skills in a cumulative fashion is one of the hallmarks of intelligence and is at the core of the Autonomous Mental Development endeavour (Weng et al. 2001). We ask our colleagues to identify the key open challenges for understanding autonomous cumulative learning in organisms and reproducing it in robots (although social processes are extremely important for human development, here we focus only on nonsocial, fully autonomous learning). As stressed by the recently funded EU Integrated Project 'IM-CLeVeR' – Intrinsically Motivated Cumulative Learning Versatile Robots (http://im-clever.eu), the problem of cumulative learning can be divided into two general sub-problems: (a) which are the signals that can drive cumulative learning? (b) which control and learning architectures can support the cumulative acquisition of skills?

Learning signals. Most research devoted to develop autonomous learning robots focuses on the solution of single tasks and hence uses task-specific learning signals. These signals have a strong limit for cumulative learning in that they can only drive the acquisition of new skills strictly related to the task(s) decided by the researcher. Inspired by the findings of both animal and human psychology, several researchers have started to explore the possibility of using non-task-specific learning signals generated by intrinsic motivations. Several fundamental issues remain open in this regard, for example: Which kind (s) of intrinsic motivations does cumulative learning need (for a useful taxonomy, see (Oudeyer, P. & Kaplan., F. 2007)? Almost all of the proposed intrinsic learning signals are based on the stimuli that the learning system perceives and processes internally, i.e. on its knowledge (e.g. Schmidhuber, J. 1991; Oudeyer et al. 2007): does cumulative learning also need intrinsic learning signals based on what the system does, i.e. on its competence (see Barto et al. 2004; Schembri et al. 2007 for some first proposals)? What are the relationships between intrinsic motivations and other motivations? What are the brain mechanisms underlying them in real organisms (e.g., phasic dopamine: Redgrave, P. & Gurney, K. 2006)?

Architectures. Cumulative learning requires that new skills are acquired on the basis of the previous ones without disrupting them. Which kind of architecture can support open-ended cumulative learning? For example, research on hierarchical reinforcement learning (Barto, A. G. & Mahadevan, S. 2003) has been developing systems based on a rather strict relationship between the functional hierarchy of skills and their structural/architectural organization. These approaches have been mainly developed at a theoretical level and for discrete problems: can they scale up to allow cumulative learning in robotic systems? How? On the other hand, it has been demonstrated that structural modularity and hierarchy are not strictly necessary for producing functionally modular and hierarchical behaviours (Botvinick, M. & Plaut, D. C. 2004; Yamashita, Y. & Tani, J. 2008). But can these approaches avoid the problem of catastrophic forgetting when requested to scale to truly open-ended cumulative learning? How? Between these two extremes there is a range of possible solutions which raise a number of open challenges: How can skills be segmented and encoded in 'soft modules'? How many hierarchical levels are needed? What are the functions of these levels and their relations? How can skills be re-used, generalised, and composed?

References:

- Weng, J., McClelland, J., Pentland, A., Sporns, O., Stockman, I., Sur, M. & Thelen, E. (2001). Autonomous mental development by robots and animals. *Science*, 291, 599–600.
- Oudeyer, P. & Kaplan, F. (2007). What is Intrinsic Motivation? A Typology of Computational Approaches. *Frontiers in Neurorobotics*, 1, E1–14.
- Schmidhuber, J. (1991). A possibility for implementing curiosity and boredom in model-building neural controllers. In Arcady-Meyer J. & Wilson S.W., (eds), *From Animals to Animats: Proceedings of the First International Conference on Simulation of Adaptive Behaviour*, Cambridge, MA : MIT Press, 222–227.
- Oudeyer, P.; Kaplan, F. & Hafner, V. V. (2007), Intrinsic Motivation Systems for Autonomous Mental Development, *IEEE Transactions on Evolutionary Computation,*11(2), 265–286.
- Barto, A., Singh, S. & Chentanez, N. (2004). Intrinsically motivated learning of hierarchical collections of skills. In Movellan, J.R., Chiba, A., Deak, G., Triesch, J., Bartlett, M. S. (eds.), *Proceedings of the Third International conference of Development and Learning*.
- Schembri, M., Mirolli, M. & Baldassarre, G. (2007). Evolving childhood's length and learning parameters in an intrinsically motivated reiforcement learning robot. In Berthouze, L., Dhristiopher, Prince G., Littman, M., Kozima, H. and Balkenius, C. (eds.), *Proceedings of the Seventh International Conference on Epigenetic Robotics*, Lund: Lund University Cognitive Studies, 141–148.

Dialog Initiation

- Redgrave, P. & Gurney, K. (2006). The short-latency dopamine signal: a role in discovering novel actions? *Nature Reviews Neuroscience*, 7(12), 967–975.
- Barto, A. G. & Mahadevan, S. (2003). Recent advances in hierarchical reinforcement learning, discrete event dynamic systems. *Discrete Event Dynamic Systems*, 13(4), 341–379.
- Botvinick, M. & Plaut, D. C. (2004), Doing without schema hierarchies: A recurrent connectionist approach to normal and impaired routine sequential action. *Psychological Review*, 111, 395–429.
- Yamashita, Y. & Tani, J. (2008). Emergence of functional hierarchy in a multiple timescale neural network model: A humanoid robot experiment. *PLoS computational biology*, 4(11), e1000220.

Call for Participation

9th IEEE International Conference on Development and Learning (ICDL 2010)

Ann Arbor, Michigan,
August 18th-21st, 2010-04-19
http://www.eecs.umich.edu/icdl-2010/

IEEE TAMD Table of Contents

IEEE TRANSACTIONS ON
AUTONOMOUS MENTAL DEVELOPMENT

Volume 1, Issue 3, October 2009   **Link:** http://ieeexplore.ieee.org/xpl/tocresult.jsp?isnumber=5402646

R-IAC: Robust Intrinsically Motivated Exploration and Active Learning

Baranes, A.; Oudeyer, P.-Y. Pages: 155-169 (pdf)

Abstract: Intelligent adaptive curiosity (IAC) was initially introduced as a developmental mechanism allowing a robot to self-organize developmental trajectories of increasing complexity without preprogramming the particular developmental stages. In this paper, we argue that IAC and other intrinsically motivated learning heuristics could be viewed as active learning algorithms that are particularly suited for learning forward models in unprepared sensorimotor spaces with large unlearnable subspaces. Then, we introduce a novel formulation of IAC, called robust intelligent adaptive curiosity (R-IAC), and show that its performances as an intrinsically motivated active learning algorithm are far superior to IAC in a complex sensorimotor space where only a small subspace is neither unlearnable nor trivial. We also show results in which the learnt forward model is reused in a control scheme. Finally, an open source accompanying software containing these algorithms as well as tools to reproduce all the experiments presented in this paper is made publicly available.

Coevolution of Role-Based Cooperation in Multiagent Systems

Yong, C. H.; Miikkulainen, R. Pages: 170-186 (pdf)

Abstract: In tasks such as pursuit and evasion, multiple agents need to coordinate their behavior to achieve a common goal. An interesting question is, how can such behavior be best evolved? A powerful approach is to control the agents with neural networks, coevolve them in separate subpopulations, and test them together in the common task. In this paper, such a method, called Multiagent Enforced SubPopulations (Multiagent ESP), is proposed and demonstrated in a prey-capture task. First, the approach is shown to be more efficient than evolving a single central controller for all agents. Second, cooperation is found to be most efficient through stigmergy, i.e., through role-based responses to the environment, rather than communication between the agents. Together these results suggest that role-based cooperation is an effective strategy in certain multiagent tasks.

What is Needed for a Robot to Acquire Grammar?
Some Underlying Primitive Mechanisms for the Synthesis of Linguistic Ability

Lyon, C; Sato, Y.; Saunders, J.; Nehaniv, C.L. Pages: 187-195 ([pdf](#))

Abstract: A robot that can communicate with humans using natural language will have to acquire a grammatical framework. This paper analyses some crucial underlying mechanisms that are needed in the construction of such a framework. The work is inspired by language acquisition in infants, but it also draws on the emergence of language in evolutionary time and in ontogenic (developmental) time. It focuses on issues arising from the use of real language with all its evolutionary baggage, in contrast to an artificial communication system, and describes approaches to addressing these issues. We can deconstruct grammar to derive underlying primitive mechanisms, including serial processing, segmentation, categorization, compositionality, and forward planning. Implementing these mechanisms are necessary preparatory steps to reconstruct a working syntactic/semantic/pragmatic processor which can handle real language. An overview is given of our own initial experiments in which a robot acquires some basic linguistic capacity via interacting with a human.

A Dynamic Systems Model of Infant Attachment **Stevens, G.T.; Jun Zhang.** Pages: 196-207 ([pdf](#))

Abstract: Attachment, or the emotional tie between an infant and its primary caregiver, has been modeled as a homeostatic process by Bowlby's (Attachment and Loss, 1969; Anxiety and Depression, 1973; Loss: Sadness and Depression, 1980). Evidence from neurophysiology has grounded such mechanism of infant attachment to the dynamic interplay between an opioid-based proximity-seeking mechanism and an NE-based arousal system that are regulated by external stimuli (interaction with primary caregiver and the environment). Here, we model such attachment mechanism and its dynamic regulation by a coupled system of ordinary differential equations. We simulated the characteristic patterns of infant behaviors in the strange situation procedure, a common instrument for assessing the quality of attachment outcomes types for infants at about one year of age. We also manipulated the parameters of our model to account for neurochemical adaptation, and to allow for caregiver style (such as responsiveness and other factors) and temperamental factor (such as reactivity and readiness in self-regulation) to be incorporated into the homeostatic regulation model of attachment dynamics. Principle component analysis revealed the characteristic regions in the parameter space that correspond to secure, anxious, and avoidant attachment typology. Implications from this kind of approach are discussed.

How Caregiver's Anticipation Shapes Infant's Vowel Through Mutual Imitation

Ishihara, H.; Yoshikawa, Y.; Miura, K.; Asada, M. Pages: 217-225 ([pdf](#))

Abstract: The mechanism of infant vowel development is a fundamental issue of human cognitive development that includes perceptual and behavioral development. This paper models the mechanism of imitation underlying caregiver–infant interaction by focusing on potential roles of the caregiver's imitation in guiding infant vowel development. Proposed imitation mechanism is constructed with two kinds of the caregiver's possible biases in mind. The first is what we call "sensorimotor magnets" in which a caregiver perceives and imitates infant vocalizations as more prototypical ones as mother-tongue vowels. The second is based on what we call "automirroring bias" by which the heard vowel is much closer to the expected vowel because of the anticipation being imitated. Computer simulation results of caregiver–infant interaction show the sensorimotor magnets help form small clusters and the automirroring bias shapes these clusters to become clearer vowels in association with the sensorimotor magnets.

A Computational Model of Acoustic Packaging

Schillingmann, L.; Wrede, B.; Rohlfing, K.J. Pages: 226-237 ([pdf](#))

Abstract: In order to learn and interact with humans, robots need to understand actions and make use of language in social interactions. The use of language for the learning of actions has been emphasized by Hirsh-Pasek and Golinkoff (MIT Press, 1996), introducing the idea of acoustic packaging . Accordingly, it has been suggested that acoustic information, typically in the form of narration, overlaps with action sequences and provides infants with a bottom-up guide to attend to relevant parts and to find structure within them. In this article, we present a computational model of the multimodal interplay of action and language in tutoring situations. For our purpose, we understand events as temporal intervals, which have to be segmented in both, the visual and the acoustic modality. Our acoustic packaging algorithm merges the segments from both modalities based on temporal overlap. First evaluation results show that acoustic packaging can provide a meaningful segmentation of action demonstration within tutoring behavior. We discuss our findings with regard to a meaningful action segmentation. Based on our future vision of acoustic packaging we point out a roadmap describing the further development of acoustic packaging and interactive scenarios it is employed in.

Developmental Stereo: Emergence of Disparity Preference in Models of the Visual Cortex

Solgi, M.; Juyang Weng. Pages: 238-252 (pdf)

Abstract: How our brains develop disparity tuned V1 and V2 cells and then integrate binocular disparity into 3-D perception of the visual world is still largely a mystery. Moreover, computational models that take into account the role of the 6-layer architecture of the laminar cortex and temporal aspects of visual stimuli are elusive for stereo. In this paper, we present cortex-inspired computational models that simulate the development of stereo receptive fields, and use developed disparity sensitive neurons to estimate binocular disparity. Not only do the results show that the use of top-down signals in the form of supervision or temporal context greatly improves the performance of the networks, but also results in biologically compatible cortical maps-the representation of disparity selectivity is grouped, and changes gradually along the cortex. To our knowledge, this work is the first neuromorphic, end-to-end model of laminar cortex that integrates temporal context to develop internal representation, and generates accurate motor actions in the challenging problem of detecting disparity in binocular natural images. The networks reach a subpixel average error in regression, and 0.90 success rate in classification, given limited resources.

Learning Posture Invariant Spatial Representations Through Temporal Correlations

Spratling, M. W. Pages: 253-263 (pdf)

Abstract: A hierarchical neural network model is used to learn, without supervision, sensory-sensory coordinate transformations like those believed to be encoded in the dorsal pathway of the cerebral cortex. The resulting representations of visual space are invariant to eye orientation, neck orientation, or posture in general. These posture invariant spatial representations are learned using the same mechanisms that have previously been proposed to operate in the cortical ventral pathway to learn object representation that are invariant to translation, scale, orientation, or viewpoint in general. This model thus suggests that the same mechanisms of learning and development operate across multiple cortical hierarchies.

Volume 2, Issue 1, March 2010   **Link:** http://ieeexplore.ieee.org/xpl/tocresult.jsp?isnumber=5429153

On the Impact of Robotics in Behavioral and Cognitive Sciences:
From Insect Navigation to Human Cognitive Development **Oudeyer, P.-Y.** Pages: 2-16 (pdf)

Abstract: The interaction of robotics with behavioral and cognitive sciences has always been tight. As often described in the literature, the living has inspired the construction of many robots. Yet, in this article, we focus on the reverse phenomenon: building robots can impact importantly the way we conceptualize behavior and cognition in animals and humans. This article presents a series of paradigmatic examples spanning from the modelling of insect navigation, the experimentation of the role of morphology to control locomotion, the development of foundational representations of the body and of the self/other distinction, the self-organization of language in robot societies, and the use of robots as therapeutic tools for children with developmental disorders. Through these examples, I review the way robots can be used as operational models confronting specific theories to reality, or can be used as proof of concepts, or as conceptual exploration tools generating new hypotheses, or used as experimental set ups to uncover particular behavioral properties in animals or humans, or even used as therapeutic tools. Finally, I discuss the fact that in spite of its role in the formation of many fundamental theories in behavioral and cognitive sciences, the use of robots is far from being accepted as a standard tool and contributions are often forgotten, leading to regular rediscoveries and slowing down cumulative progress. The article concludes by highlighting the high priority of further historical and epistemological work.

Neuromorphically Inspired Appraisal-Based Decision Making in a Cognitive Robot

Gordon, S. M.; Kawamura, K.; Wilkes, D. M. Pages: 17-39 (pdf)

Abstract: Real-time search techniques have been used extensively in the areas of task planning and decision making. In order to be effective, however, these techniques require task-specific domain knowledge in the form of heuristic or utility functions. These functions can either be embedded by the programmer, or learned by the system over time. Unfortunately, many of the reinforcement learning techniques that might be used to acquire this knowledge generally demand static feature vector representations defined a priori. Current neurobiological research offers key insights into how the cognitive processing of experience may be used to alleviate dependence on preprogrammed heuristic functions, as well as on static feature representations. Research also suggests that internal appraisals are influenced by such processing and that these appraisals integrate with the cognitive decision-making process, providing a range of useful and adaptive control signals that focus, inform, and mediate deliberation. This paper describes a neuromorphically inspired approach for cognitively processing experience in order to: 1) abstract state information; 2) learn utility functions over this state abstraction; and 3) learn to tradeoff between performance and deliberation time.

Reproducing Interaction Contingency Toward Open-Ended Development of Social Actions: Case Study on Joint Attention **Sumioka, H.; Yoshikawa, Y.; Asada, M.** Pages: 40-50 ([pdf])

Abstract: How can human infants gradually socialize through interaction with their caregivers? This paper presents a learning mechanism that incrementally acquires social actions by finding and reproducing the contingency in interaction with a caregiver. A contingency measure based on transfer entropy is used to select the appropriate pairs of variables to be associated to acquire social actions from the set of all possible pairs. Joint attention behavior is tested to examine the development of social actions caused by responding to changes in caregiver behavior due to reproducing the found contingency. The results of computer simulations of human–robot interaction indicate that a robot acquires a series of actions related to joint attention such as gaze following and alternation in an order that almost matches the infant development of joint attention found in developmental psychology. The difference in the order between them is discussed based on the analysis of robot behavior, and then future issues are given.

Computational Developmental Neuroscience: Capturing Developmental Trajectories From Genes to Cognition **Thivierge, J.-P.** Pages: 51-58 ([pdf])

Abstract: Over the course of development, the central nervous system grows into a complex set of structures that ultimately controls our experiences and interactions with the world. To understand brain development, researchers must disentangle the contributions of genes, neural activity, synaptic plasticity, and intrinsic noise in guiding the growth of axons between brain regions. Here, we examine how computer simulations can shed light on neural development, making headway towards systems that self-organize into fully autonomous models of the brain. We argue that these simulations should focus on the "open-ended" nature of development, rather than a set of deterministic outcomes.