IEEE CIS

AMD Newsletter

Robotics
Psychology
Neuroscience
Machine Intelligence

The Newsletter of the Autonomous Mental Development Technical Committee

VOLUME 6, NUMBER 2 ISSN 1550-1914 October 2009

Editorial

This special issue features a dialog of essential importance for our field: Can we identify principles that specifically define developmental robotics? And should we strive to organize research around such principles or rather consider tinkering and ad hoc investigations as a strength or a marker of scientific innovation as argued by philosopher Feyerabend? Coordinated by Max Lungarella, we see compelling arguments for both conceptions put forward by Josh Bongard, Franck Guerin, Chrystopher Nehaniv, Linda Smith, Alex Stoychev, Juyang Weng, and Patricia Zukow-Goldring.

Then, a novel call for dialog is proposed by Angelo Cangelosi and relates to the central and controversial problem of language development and grounding: « The Symbol Grounding Problem Has Been Solved: Or Maybe Not? ». Interested researchers are welcome to submit a response (contact a.cangelosi@plymouth.ac.uk or pierre-yves.oudeyer@inria.fr) by February 28th, 2009. The length of each response must be between 300 and 500 words (including references).

Finally, I encourage strongly submissions to the ***special issue of the IEEE TAMD*** journal ***on active learning and intrinsically motivated exploration in robots***, guest edited by Manuel Lopes and myself. The call for paper can be found at: http://flowers.inria.fr/tamd-activeLearningIntrinsicMotivation.htm (deadline for submission: January 31st, 2010). This special issue is jointly supported by the IEEE CIS Technical Committee on AMD and by the IEEE RAS Technical Committee on robot learning.

-Pierre-Yves Oudeyer, INRIA, Editor

AMD TC Chairman's Message

After the successful launch of the inaugural issue of TAMD, we have now published the second issue. As for the inaugural issue, we are publishing the table of contents in this newsletter. I want to thank all the authors who have submitted papers, all the Associate Editors for managing the review process, and all the reviewers for volunteering their precious time. We are now working hard for the remaining two issues for this year. For next year, besides regular submissions, we will have two special issues: (1) active learning and intrinsic motivation; and (2) cognitive vision (tentative). We look forward to receiving more submissions for the forthcoming issues as well as for the special issues.

In terms of conference organization, AMD TC members have successfully organized ICDL 2009 and the 2009 IROS Workshop on AMD, and co-organized Epirob 2009. Please read the conference report.

The time has never been better as well as critical for our AMD community and its growth. The importance of AMD methodologies is being increasingly recognized by the more traditional computational intelligence and robotics communities. Your active participation in attending ICDL and EpiRob, in organizing AMD special sessions and symposia in relevant conferences, and in contributing papers to and reviewing papers for IEEE TAMD, is vital to this AMD community. I want you!

-Zhengyou Zhang, Current chair of the AMD TC

Committee News

- ICDL 2009 was successfully held in Shanghai, China, June 4-7, 2009. See conference report.
- EpiRob 2009 is being held in Venice, Italy, November 12-14, 2009. You are invited to attend.
- Danil Prokhorov has successfully organized the *IROS 2009 workshop on Autonomous Mental Development for Intelligent Robots & Systems* on October 11, 2009. It was well attended with about 30 people.
- ICDL 2010 will take place in Ann Arbor, Michigan, USA, August 18-21, 2010. Call for Papers will be distributed shortly.
- The IJCNN 2010 conference (part of IEEE WCCI 2010) will take place in Barcelona, Spain, July 18-23, 2009. More info at http://www.wcci2010.org/special-sessions. Special sessions on AMD related topics are solicited. Please contact Vincenzo Piuri at vincenzo.piuri@unimi.it and Zhengyou Zhang at zhang@microsoft.com.
- IEEE SSCI 2011 will take place in Paris in April 2011. Anyone who is interested in organizing a symposium on AMD please contact Zhengyou Zhang at zhang@microsoft.com.

Dialog Column

Developmental Robotics: From Black Art to Discipline Guided by Principles?

Max Lungarella
Artificial Intelligence Laboratory, University of Zurich, Switzerland

Developmental/epigenetic robotics differs from traditional robotics (and artificial intelligence) in at least three crucial aspects:

1. There is a strong emphasis on body structure and environment as causal elements in the emergence of organized behavior and cognition, requiring their explicit inclusion in models of emergence and development of cognition;
2. Artificial cognitive systems are not simply programmed to solve a specific task, but rather a developmental process is initiated and maintained during which cognition emerges and develops through a process of self-organization and co-development (and interaction) between the artificial organism and its surrounding environment;
3. In contrast to robotics as well as traditional disciplines such as physics, and mathematics which are described by basic axioms and fundamental laws, the basic principles governing the dynamics of artificial (and natural) developmental systems are still largely unknown;

My question for the readership of this dialog column is: Are there any laws governing developmental systems or even a theory, and if so, how can such laws be turned into design principles for engineering robots which are more autonomous, adaptive, or resilient? Or, more fundamentally, is it plausible to assume that an approach to the construction of intelligent autonomous systems guided by principles is preferable to one which relies on ad hoc mechanisms? On the one hand, it could be argued that such principles have essentially two advantages: 1) they allow capturing design ideas and heuristics in a concise and pertinent way, and 2) they reduce the amount of tinkering and blind trial-and-error. On the other hand, it could also be reasoned that biological evolution itself is based on tinkering and blind trial-and-error, and yet has produced extremely adaptive creatures – implying that ad hoc mechanisms might actually work after all, if given sufficient time and raw materials.

To start the discussion, here I present three candidate design principles (see also Stoychev; 2006; Pfeifer et al., 2007; Smith and Breazeal, 2007):

1. When designing a developmental agent it is important to see the behavior of the agent not merely as the outcome of an internal control structure (such as the central nervous system). A system's behavior is also affected by the ecological niche in which the system is physically embedded, by its morphology (the shape of its body and limbs, as well as the type and placement of sensors and effectors), and the material properties of the elements composing the morphology;
2. An embodied agent does not passively absorb information from its surrounding environment: coupled sensory-motor activity and body morphology induce statistical regularities in the sensory input as well as in the control architecture and therefore enhance internal information processing and therefore learning. This property should thus be taken into account at design time;

Dialog Column

3. Viewing an embodied agent as a complex dynamical system enables us to employ concepts such as self-organization and emergence rather than hierarchical top-down control. In other words, the agent should not be completely hardwired at the outset, but the agent needs to be endowed with the ability to self-direct the exploration of its own sensory-motor capabilities and with means to escape its limited built-in behavioral repertoire, and to acquire its own history;

Clearly, a large number of such design principles can be abstracted from biological systems, and their inspiration can take place at several levels, ranging from a "faithful" replication of biological mechanisms to a rather generic implementation of biological principles leaving room for dynamics intrinsic to artifacts but not found in natural systems. But then, how does one choose their level of generality? Will it eventually be possible to turn developmental robotics from a black art into a principled discipline?

References:

- Pfeifer, R., Lungarella, M. and Iida, F. (2007). Self-organization, embodiment, and biologically inspired robotics. Science, 318:1088–1093.
- Smith, L.B. and Breazeal, C. (2007). The dynamic lift of developmental process. Developmental Science, 10(1):61–68.
- Stoytchev, S. (2006). Five basic principles of developmental robotics. NIPS Workshop on Grounding, Perception, Knowledge, and Cognition.

We Need a Success, Then We Can Generalise Principles

Franck Guerin
Department of Computing Science, University of Aberdeen, Scotland

In the nineteenth century heavier than air flying machines were not getting off the ground. The only examples of flying machines available were examples of failure. At that time anyone who proposed principles could not be taken too seriously, as there was no example of success. It is hard to generalise from a failure. A failure shows one way of doing it wrongly. There were some examples of attempting to generalise from failure, such as Lord Kelvin's "Heavier-than-air flying machines are impossible" (Kelvin, 1895). Any principle generalised from some failures is generalising over a huge number of cases for which there is no experimental evidence. A success on the other hand can lead to generalisations over a much narrower set of cases, with more evidence to support them. A success can be repeated and experimented with to see the range over which various parameters can be varied, and how this impacts the performance. It is then possible to abstract some principles relating design decisions to outcomes.

With developmental robotics I feel we have not got off the ground yet. Getting off the ground here would mean building a system which can display ongoing development, continuously building new knowledge on top of what it knows, like a human infant. We have some systems that develop a little and then stop. We are nowhere near having a convincing example of ongoing development. Prince et al.'s survey comes to the same conclusion (Prince et al., 2005): "This leads us to view current examples of epigenetic robots as demonstrating emergence, but not ongoing emergence." So long as we have only got examples of failures it is premature to generalise principles, even damaging. Generalising from successful biological systems is also suspect because we do not know which aspects of those systems are essential to their success (just as early flight researchers did not know if feathers were essential on their strap-on wings).

A principle is a constraint, and advises against exploring certain regions of the design space. Given that we currently don't know how to build developing robots it does not seem wise to rule out the exploration of any avenues. Consider the principle that AI work must be done on real robots in the real world, which are complete intelligent systems at each step. This advises against working on isolated aspects of the problem in a simplified simulation. The principle has been highly influential, but not everybody agrees:
"The worst fad has been these stupid little robots," said Minsky. "Graduate students are wasting 3 years of their lives soldering and repairing robots, instead of making them smart. It's really shocking." (Baard, 2003).

It is noteworthy that the Wright brothers made great progress when they isolated part of the problem by building a wind tunnel and studying the lift on various wing shapes. It would seem that if there is a principle developmental robotics should follow it is not yet a design principle for engineering robots, but rather the good old scientific principles: observe, hypothesise and test.

References:

- Thomson, W. a.k.a. Lord Kelvin (1895), as quoted in The Experts Speak: The Definitive Compendium of Authoritative Misinformation (1984) by Christopher Cerf and Victor Navasky, p. 236
- Prince, C., Helder, N., and Hollich, G. (2005). Ongoing emergence: A core concept in epigenetic robotics. In Berthouze, L., Kaplan, F., Kozima, H., Yano, H., Konczak, J., Metta, G., Nadel,J., Sandini, G., Stojanov, G., and Balkenius, C., editors, Proceedings of EpiRob05 - International Conference on Epigenetic Robotics, pages 63–70. Lund University Cognitive Studies.
- Baard, M. (2003) AI Founder Blasts Modern Research. In Wired Magazine http://www.wired.com/science/discoveries/news/2003/05/58714

AMD Principles: Have We Passed "Black Art"?

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

Max Lungarella has used the term "black art" to refer to the developmental robotics field, which I understand is truly the feeling among many roboticists. In my humble opinion, however, "black art" does not seem to fit the current status of AMD since the AMD field includes natural intelligence disciplines as it should be, at least if we take what the Workshop on Development and Learning 2000 (1st ICDL, http://www.cse.msu.edu/dl/) discussed. There have been many principles studied in or that resulted from AMD/developmental robotics studies, especially *qualitative* design principles that distinguish AMD from traditional machine learning in robotics. *The task-nonspecificity* principle Max mentioned was such a qualitative principle. It was published as early as in Jan. 2001 along with a few other major ones. I do not mean that these principles are sufficient. For example, the use of developmental motivation systems later published is also necessary for effective AMD. However, "black art" does not seem to be a reasonable and fair assessment of the current work in developmental robotics. I hope that Max probably did not mean it either.

As a relatively new area in robotics, developmental robotics should first identify *qualitative* principles that clearly *distinguish* AMD and developmental robotics from *traditional machine learning and robotics*, such as many Bayesian learning methods that have a long history and beautiful theories. Why can these Bayesian learning methods not enable autonomous mental development, or at least have not yet? Consider their handcrafted task-specific representations. Once a task-specific representation is handcrafted into a robot, how is the robot limited? Can it learn tasks not modeled by handcrafting?
None of the three design principles Max raised seem to be *unique* to *developmental* robotics. The first principle about the use of *body morphology* is well known in traditional robotics --- even *inverse kinematics* use it extensively. The second principle is called *active learning* in psychology, AI and traditional machine learning in robotics (e.g., SLAM and POMDP use it). The third one on *embodiment* is taken for granted in traditional robotics (i.e., any robot has a body to sense and act), probably not in GOFAI. *Self-organization* and *emergence* are great points, but they are well known in neural networks for along time --- e.g., SOM is well known and almost all existing neural networks handle emergence. By the way, I humbly raise a point that, interestingly, self-organization and emergence are *not inconsistent* with "hierarchical top-down control." Top-down control is well documented in neuroscience, e.g., top-down attention control - see a classical review article by (Desimone & Duncun, 1995). It has at least two types --- *position based* and *object* (or *feature*) *based*. That is, top-down control is also emergent through development.

Let us not repeat well-known principles in non-developmental robotics as the design principles of *developmental robotics* so that the term "developmental robotics" has some well identifiable credibility. Max's three principles, body morphology, activeness, and embodiment, are all good and valid points. But they are not new to non-developmental robotics and not unique to AMD either.

Further, there is a lack of *computational* design principles that are suited for AMD but also distinguish AMD from tradi-

Dialog Column

tional machine learning. For example, there is a scarcity of computational models for top-down attention control and its task-nonspecific developmental mechanisms, although there is no lack of *qualitative* principles for top-down attention control (Desimone and Duncan, 1995; Buschman and Miller, 2007).   My coworkers and I have only started a computational model of (developmental) top-down attention (Ji et al., 2008). In the model, different features do take a hierarchy, but such a hierarchy is dynamic, transient, and context-dependent. Along this line, there are many interesting open problems waiting for us to solve.

How to correct this "black art" feeling among many roboticists is also an important topic. The lack of intellectual communication between this AMD community and the larger robotics community and other related communities is contributing to this false "black art" feeling. Let us continue to discuss how to correct this unfortunate situation.

References:

- Desimone, R. and Duncan, J. (1995) Neural mechanisms of selective visual attention, *Annual Review of Neuroscience*, vol. 18, pp. 193-222.
- Buschman, T. J. and Miller, E. K. (2007) Top-down versus bottom-up control of attention in the prefrontal and posterior parietal cortices, *Science*, vol. 315, pp. 1860-1862.
- Ji, Z., Weng, J., and Prokhorov, D. (2008) Where-What Network 1: 'Where' and 'What' Assist Each Other Through Top-down Connections, in *Proc. IEEE International Conference on Development and Learning*, Monterey, CA, pp. 61-66, Aug. 9-12.

Developmental Robotics: Black Art or a Discipline Guided by Principles?

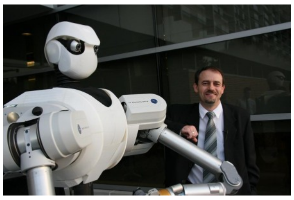

Alexander Stoytchev
Developmental Robotics Laboratory
Iowa State University, U.S.A.

Is Developmental Robotics a discipline guided by principles or is it a black art? I was struggling with this same question three years ago when I wrote (Stoytchev, 2006). In that paper I formulated five basic principles of Developmental Robotics, which can be used to guide research in the field until the principles are either proven wrong or replaced with new ones. So far, these principles have served my lab's research very well.

It was very hard to write that paper for several reasons. First, I had no good starting points as nobody at that time had attempted to do something similar. Second, I wanted to avoid the pitfall of listing a set of unrelated principles that may conflict with each other. Finally, I wanted to derive the principles from the latest research in the field (and related disciplines) so that they are not detached from reality. I managed to balance these constraints by identifying one main principle and logically deriving all others from it. This main principle is the verification principle and it is the subject of this column.
The **verification principle** states that: "An AI system can create and maintain knowledge only to the extent that it can verify that knowledge itself." (Sutton, 2001)

This principle recognizes and directly addresses the most common scalability challenge in AI and robotics, namely, the numerous hidden assumptions made by the programmers of these systems. Some of these assumptions are unnecessary. Others are valid only for a narrow domain. Sooner or later, however, the hidden assumptions are violated and the robot is left with few good options to recover from these situations. Thus, scalability and reliability become next to impossible. The principle is easy to state. Embracing it, however, changes almost everything. It changes how the robot extracts useful information (it uses its behavioral repertoire). It changes how the robot represents learnable quantities (in terms of its sensorimotor repertoire and relative to its own body). It changes the order in which the robot explores the environment (from the most verifiable to the least-verifiable). See (Stoytchev, 2006) for more details.

In summary, making the commitment to follow the verification principle is a real game changer. A robot that is programmed with this principle in mind can autonomously test, verify, and correct its own knowledge representation without requiring the intervention of a human programmer. Embracing this principle would put Developmental Robotics on a more stable footing.

Dialog Column

Acknowledgements:
Many thanks to Max Lungarella for initiating this dialog. This is exactly the question that the field as a whole should be addressing. Maybe we should organize a workshop on this topic.

References:

- Stoytchev, A. (2006) "Five Basic Principles of Developmental Robotics", NIPS 2006 Workshop on Grounding Perception, Knowledge and Cognition in Sensori-Motor Experience, Whistler, British Columbia, Canada.
- Sutton, R. (2001) "Verification, the key to AI," on-line essay http://www.cs.ualberta.ca/~sutton/IncIdeas/KeytoAI.html

---

A Riddle Wrapped in a Mystery

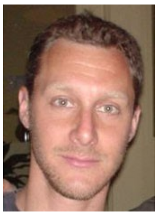

Josh Bongard
Morphology, Evolution and Cognition Laboratory
University of Vermont, USA

Lungarella points up the fundamental divide between traditional robotics and the loose confederation of bio-, developmental, and evolutionary robotics and artificial life. The former camp emphasizes hand-designing behavior, while the latter stresses the establishment of longer-term processes which in turn give rise to useful behavior. The appeal of the latter approach is indeed, as Lungarella describes, that it seems to hold more promise for delivering complex autonomous machines as our ability to directly program complex systems is limited. This advantage is paid for however with the loss of formality: our field(s) are at the state where we cling to ad hoc design rules for establishing learning and/or evolutionary processes which sometimes produce useful behaviors.

In my own home field of evolutionary robotics (Nolfi and Floreano, 2004) this ad hoc nature can be seen in the wild profusion of approaches. Some attempt to evolve control policies while other co-evolve both morphology and control (Sims, 1994). Some fold development into the evolutionary process such that rather than evolving the control and/or morphology of a robot directly, the evolutionary process evolves a growth program which then guides the development of the robot (Bongard, 2002). Across these approaches vastly different evolutionary algorithms are employed, and evolved robots are tested in simulation, constructed and tested in reality, or both (Bongard et al., 2006).

Indeed evolutionary robotics can be seen as an even 'darker art' than developmental robotics. Certain aspects of natural selection are incorporated into the evolutionary algorithm while others are abstracted away, and these processes in turn guide developmental processes which themselves are relatively arbitrary simulations of biological development. However, many see this very ad hoc nature of the field as it now stands as a strength (Pfeifer and Bongard, 2007). Many of the turning points in nascent scientific disciplines arose from chaotic and ad hoc investigation (Kuhn, 1970). Because of this, it has been argued that scientists should proceed counter-inductively (Fereyabend, 1993): we should work to generate evidence that degrades current theories and thereby pave the way for new ones. I hold that bio-, developmental and evolutionary robotics as well as artificial life is indeed working to do just this: the theory that intelligence can be programmed directly into machines by humans seems increasingly implausible; we should instead work to show that evolution and development can give rise to intelligent entities not only in carbon, but in silicon as well.

References

- Nolfi, S., Floreano, D. (2004). *Evolutionary Robotics: The Biology, Intelligence and Technology of Self-Organizing Machines*. Bradford Books.
- Sims, K. (1994). Evolving 3D morphology and behavior by competition. *Procs. of the Fourth Intl. Workshop on the Synthesis and Simulation of Living Systems*, pp. 28-39.
- Bongard, J. (2002) Evolving modular genetic regulatory networks. *Procs.of the Congress on Evolutionary Comptuation*, pp. 17-21.
- Bongard, J., Zykov, V., Lipson, H. (2006). Resilient machines through continuous self-modeling. *Science* 314(5802): 1118-1121.
- Pfeifer, R., Bongard, J. (2007). *How the Body Shapes the Way We Think: A New View of Intelligence*. MIT Press.
- Kuhn, T. S. (1970). *The Structure of Scientific Revolutions*. University of Chicago Press.
- Feyerabend, P. K. (1993). *Against Method*. Verso Books.

Dialog Column

Developing Intelligence

Linda B. Smith
Cognitive Development Lab, Indiana University, USA.

People, at least grown up people, are smart, and so they are the benchmark against which intelligent systems (or animals) are compared. But people do not start off obviously smart. Indeed, human babies can do very little: they cannot control movements nor can they self-regulate arousal and bodily functions. But they become smart. How does one explain the "specialness" of human intelligence? How can one leverage this smartness to build smart machines? Evolution provides an important clue. Many species do very smart things –navigate, calculate, detect and smartly use subtle forms of regularities. But in many species, this intelligence is limited to specific tasks and contexts and is not transportable, nor inventive. In contrast, the species (humans, great apes) we think of as having the most advanced forms of biological intelligence are advanced precisely because they are open systems, influenced by many sources of information, generalizing broadly and inventing new solutions. So, here are the design principles from human development:

- Developmental process (like evolution) is local, individualistic, and opportunistic and highly creative;
- Many heterogeneous processes (e.g., seeing and hearing, moving and seeing) are timelocked to each other and to the world and as a product of their couplings, change their internal dynamics (and thus the computations they can perform);
- Many different tasks create different overlapping soft assemblies of these different component processes and thus cascading effects of learning from one task to others;
- There are many different overlapping soft-assemblies of these different components such that overlapping integrations in many different tasks create abstract processes that transcend the specifics of specific modalities and the here-and-now of specific tasks;
- Closed loops such that each action creates a new opportunity for perception and learning and such that the learner can perceive the consequences of her own actions, the images that the learner creates in the world.

The *developmental process* is complex, multi-leveled, multi-component and in its entirety, it makes human intelligence what it is, and what it *can be*.

From Being a Body to Becoming an Intelligent Agent

Patricia Zukow-Goldring
Center for the study of women,
University of California, USA

How does a novice come to know what everybody else already knows? But first, what exactly do those competent or adept others know? They can recognize, participate in/understand, and communicate about ongoing events. These events entail a structure and organization that is emergent depending on its quiddity or particulars of just who, what, when, and where (Garfinkel, 1967). An ecological niche or the environment, whether "natural" or "laboratory", is cultural. However, there is ***no*** parity. Why? Because each intelligent agent has a different "architecture", achieves a different developmental level or history of experience, at any given moment moves along a different path, and thus occupies a different observing point. Due to these differences, individuals negotiate a common understanding for all practical purposes.

Gibson (1979) proposed the notion of affordances referring to the ability of creatures to perceive opportunities for action in their environment which underscores the inextricable coupling of perceiving and acting. Others added the notion of effectivities—the repertoire of what the body can do—as a dual complement of affordances. The infant's discovery of a range of affordances and effectivities contributes to participating in a new activity. Especially relevant to this idea is the young infant's known ability to detect regularities or invariants in the continuous stream of perceptual information. However, objects do not tell us what to do. Someone shows us how to use them. Further, contrary to the assumptions of many, spontaneous imitation of something new rarely occurs (Zukow-Goldring, 2008). Instead, caregivers educate attention by bracketing ongoing actions with gestures that direct the infant's attention to perceptual information embodied in action sequences. Such supervised learning narrows the search space and enhances the speed of engaging adeptly in a new activity and provides a basis for achieving a common understanding of ongoing events.

Dialog Column

When caregivers embody their infants (put them through the motions), the infants unavoidably pick up proprioceptive information that tells them what their bodies can do, specifying the body's work (posture, effectivities/movements). As infants move in new ways, they can detect previously unavailable perceptual information (in vision, touch, taste, smell, sound) that specifies what other things afford for action. Detecting the synchronous onset/offset, tempo, rhythm, and intensity across modalities may point both ways when learning new actions: to the emergent, inextricable linking of effectivity and affordance for the self and provides opportunities to notice that the self is "like the other." The synchrony across different sensory modalities may "bind" the different cues together. Without this binding, it may not be possible to assess what are the relevant cues and what is "background noise." These caregiver practices during assisted imitation may illuminate how automata might detect and learn new affordances for action by observing and interacting with other intelligent agents.

References:

- Garfinkel, H. (1967). Studies in ethnomethodology. Englewood Cliffs, NJ: Prentice Hall.
- Gibson, J. J., 1979, The Ecological Approach to Visual Perception, Boston: Houghton Mifflin.
- Zukow-Goldring, P. (2009). Assisted imitation: Caregiver gestures cultivate self-other equivalence, contributing to a shared understanding of events. Multimod Conference.

Ontogenic Robots, Mechanisms and Meaningful Development

Chrystopher L. Nehaniv
Adaptive Systems Research Group,
University of Hertfordshire, U.K.

While it is possible to model in various ways developmental processes using simulation, hardware, and theories of ontogeny extracted from the literature, much of this work has been comprised of ad hoc exercises in putting together demo engineering simulations. Unfortunately neuroscientists and psychologists are often not in a position able to assess to what extent fundamental properties of development have been captured (if any), and how much of what they are seeing is merely kludged together bits and pieces sprinkled with some minimal algorithmic glue. This often appears to researchers who build such systems to be 'smoke and mirrors' or 'just cheating' to varying extents.

On the other hand, Rodney Brooks speaking of life and intelligence has said, at least in informal conversation, that "it's all just cheating". Indeed, the evolution of life and the development of naturally intelligent systems takes anything it can use. Fair enough. No one is prohibiting life from 'cheating'. But, as this dialog is intended to expose, are there organizational principles that can generally be applied in developmental robotics and that also apply in living organisms that grow and develop? (1)

Growing evidence suggests that the answer seems to be 'yes', where the mechanisms or classes of mechanisms include harnessing each of phenotypic plasticity (West-Eberhard 2003), evolvability within cells, embryos and organisms (Kirschner & Gerhart, 1997, 1998), freezing and releasing degrees of freedom (Bernstein, 1967), morphological computing (Pfeifer and Bongard 2006), as well as harnessing social interaction and primary intersubjectivity to expand zones of proximal development and temporal extended experience (Vygotsky 1978, Trevarthen 1979, Nehaniv and Dautenhahn 2003, 2005, Mirza et al, 2006) with related methods that balance curiosity and mastery of skills (Steels 2004, Csikszentmhaly 1991, Kaplan & Oudeyer 2006).

Self-organizing processes expressible in terms of Information Theory (in the sense of Claude Shannon) seem to play an underlying role in understanding development and its relationship to evolution and learning (e.g., Nehaniv et al. 2007). A close coupling of perception and action embedded in a context of embodiment clearly plays a key role (Varela, Thompson, Rosch 1991) and useful any notion of (generally maligned and ignored) representation must emerge from this enactive level (Millikan 2005, Klyubin et al 2008).

One process that we are quite sure plays a role at multiple levels is Darwinian evolution. Actually, this is the only mechanisms we know for creating purpose in material systems (cf. Dennett, 1996).

Dialog Column

Meaningful information in sensors, actuation, and internal 'mental development' in interaction with the physical and social environment only gains meaning from the perspective of the agent itself when agent's structure reflects the evolutionary history of a lineage subject to the constraints of evolution (see Nehaniv 2005 for an axiomatization of Darwinian evolution and Nehaniv, Dautenhahn, & Loomes, 1999, for meaning for observers and agents). As far as science knows, it is only in the setting of such an enactive fully embodied process that any of the above mentioned mechanisms acquires self-generated goal-directedness . Perhaps only when our ontogenic robots that grow up do arise as individuals in evolving populations -- or perhaps as individuals that are essentially copies of such individuals, e.g. a nano-level copy of a particular human-being -- then we will capture the essential principles that give meaning and purpose to ontogeny.

References and footnotes are available at http://flowers.inria.fr/AMDREF-v6n2-09.pdf

---

Reply and Summary:
Developmental Robotics: From Black Art to Discipline Guided by Principles?

Max Lungarella
Artificial Intelligence Laboratory, University of Zurich, Switzerland

Let me start by thanking my colleagues for their stimulating replies to the dialog initiated in April, where I asked the following questions: First, is an approach to the construction of intelligent autonomous systems guided by principles preferable to one which relies on ad hoc mechanisms? Second, are there any laws governing developmental systems or even a theory? And third, how can such laws, if they exist, be turned into design principles for engineering autonomous, adaptive, and resilient machines?

The responses to the first question largely fall in two camps. On the on hand, Guerin writes that “a principle is a constraint, which advises against exploring certain regions of the design space” and as such in a field like developmental robotics which “has not yet got off the ground can be even damaging.” Similarly, Bongard (extending his argument to include also evolutionary robotics) states that “many see the very ad hoc nature of the evo-devo field as it now stands as a strength”, and paraphrasing Kuhn he adds “that many of the turning points in nascent scientific disciplines arose from chaotic and ad hoc investigation” (what I called a state of a black art and magic).

On the other hand, the contributions of Weng, Stoytchev and Nehaniv seem to indicate that developmental robotics is not in a phase of chaotic and ad hoc investigation (anymore). Weng states that “developmental robotics has a series of qualitative design principles” and Stoytchev even presents one (the verification principle), while Nehaniv is more careful and writes that growing evidence points to “organizational principles that can be generally applied to developmental robotics.”

Now, both camps make good points in favor or disfavor of their arguments, and it is probably the case that only the future will tell. In favor of an approach relying on principles, I can think of at least one methodological advantage in the context of developmental robotics: Many researchers make theoretical assumptions about processes and mechanisms whenever they design an experimental system to conduct a study (be it a robot or a computer program). Although such assumptions are typically implicit, they often strongly influence all subsequent decisions and outcomes. Design principles allow making such implicit assumptions explicit and precise. Operationally, one way of implementing this idea is to carefully *observe* a developmental system and to *hypothesize* general principles of adaptive behavior based on the assumption that some of those principles might be at work in other systems or at other levels as well.

Experimental scenarios are then devised to quantify and *test* the proposed principles. It follows that seeking design principles is not necessarily in contradiction with the common scientific methodology of observing, hypothesizing and testing (as argued by Guerin). The testing is done, by the way, with robots, which in this case are merely scientific tools of investigation, not any different from conventional computational models.

Dialog Column

The question is what insights gained from the "careful observation" of developmental systems could and should be exploited for designing robots. Simply copying a biological system is either not feasible (even a single neuron is too complicated to be synthesized artificially in every detail), is of little interest (animals must satisfy multiple constraints that do not apply to robots, such as keeping their metabolism up and running), or the technological solution is superior to the one found in nature with respect to some measurable criteria. Rather the goal is to work out principles of developmental systems and transfer those to robot design. This philosophy underlies, for instance, the rapidly expanding field of bionics, which seeks to design technology by mimicking the salient features of biological structures (Vincent et al., 2006). Smith and Zukow-Goldring, for instance, propose a set of design principles from human development. The question is how can one map these principles onto design principles for robots? Of course, more thought will have to be devoted to translate their design principles into guidelines for engineering developmental robots. It is indeed a hairy issue, and I am afraid that more discussion will be necessary to reach agreement.

Before concluding this response, I would like to dispel some common misunderstandings about embodiment, and that is, the reciprocal and dynamic coupling of brain (control), body, and environment. Weng states that the use of "body morphology" is well known in traditional robotics. That is indeed the case, but does traditional robotics exploit body morphology? To my knowledge, most of the robots in use are built using the very same materials, and are controlled using the very same algorithms. Traditional control engineering approaches strive to avoid, or actively suppress, nonlinear dynamic coupling among components, and that is, body morphology. Biological organisms, by contrast, have evolved to perform and survive in a world characterized by rapid changes, high uncertainty, indefinite richness, and limited availability of information. The body is typically reduced to a set of typically immutable equations and the environment is highly controlled with no or very little uncertainty. The embodied approach suggests using the body morphology as an additional design parameter. I refer to Pfeifer et al. (2007) for many examples where such use has been successfully employed. To conclude this dialog, my questions remain largely unanswered. Only the future will tell what approach is the right one.

References:

- Vincent, J.V.F., Bogatyreva, O.A., Bogatyrev, N.R., Bowyer, A., Pahl, A.-K. (2006). Biomimetics: its practice and theory. *J. R. Soc. Interface*, 3: 471–482.
- Pfeifer, R., Lungarella, M. and Iida, F. (2007). Self-organization, embodiment, and biologically inspired robotics. *Science*, 318: 1088–1093.

Dialog Initiation

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-

Dialog Initiation

symbol link is the main property that differentiates a real symbol from an index, as in Peirce's semiotics.

These symbolic links also supports 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 system 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. 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 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 understading of human cognition, of should all models of language learning be based solely on grounding mechanisms?
4. Does cognitive development really plays 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 very easily language. 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

Dialog Initiation

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 are available at http://flowers.inria.fr/AMDREF-v6n2-09.pdf

Conference Reports

8th International Conference on Development and Learning (ICDL-2009)

Jochen Triesch and Zhengyou Zhang

The IEEE 8th International Conference on Development and Learning (ICDL) was held in Shanghai, China, during June 4-7, 2009. Like in previous years, this ICDL offered an exciting single track program of contributed talks. The program was complemented by two poster sessions facilitating in-depth discussions, and preceded by two tutorials by Tiande Shou and Juyang Weng. Last but not least, we had a wonderful set of invited speakers including Mitsuo Kawato, Andrew Parker, Mriganka Sur, and Manabu Tanifuji in the program.

Like for last year's ICDL, we had allowed for two types of submissions: regular papers and one-page abstracts for "late-breaking" results. In addition, Giorgio Metta, Gordon Cheng and Tamim Asfour organized a special session. We received 71 submissions in total - a small number for ICDL. These submissions were reviewed with the help of 16 area chairs, 63 program committee members and 3 additional reviewers. Apart from a few exceptions, each full paper submission received at least three reviews. The quality of the submissions was generally high, so that 55 made it into the final program.
Many people helped with the organization of the conference. Our organizing committee included Juyang Weng (general chair), Tiande Shou and Xiangyang Xue (general co-chairs), Jochen Triesch and Zhengyou Zhang (program chairs), Yilu Zhang (publication chair), Alexander Stoytchev (publicity chair), Hiroaki Wagatsuma, Pierre-Yves Oudeyer, and Gedeon Deák (publicity co-chairs), and Hong Lu (local organization chair) and her local organization team.

Financial sponsoring was obtained by the IEEE Computational Intelligence Society and technical sponsorship by the Cognitive Science Society. Microsoft Research and the Cognitive Science Society sponsored best paper awards and student travel awards. Last but not least, Fudan University sponsored a memorable boat trip on the Huangpu River from which we enjoyed fantastic views of Shanghai. Note that the next ICDL will be held during August 18-21, 2010 in Ann Arbor, Michigan, organized by Benjamin Kuipers and Thomas Shultz (general chairs) and Alexander Stoytchev and Chen Yu (program chairs). We look forward to seeing you there!

Publication

IEEE Transactions on Autonomous Mental Development (TAMD)

Table of Contents: Volume 1, Number 2, August 2009

http://ieeexplore.ieee.org/xpl/RecentIssue.jsp?punumber=4563672

Homepage with submission instructions: http://ieee-cis.org/pubs/tamd/
Submission website: http://mc.manuscriptcentral.com/tamd-ieee