IEEE CIS

AMD Newsletter

Robotics
Psychology
Neuroscience
Machine Intelligence

The Newsletter of the Autonomous Mental Development Technical Committee

VOLUME 5, NUMBER 1 ISSN 1550-1914 April 2008

Message from the editor

After five years of existence, the IEEE CIS AMD Newsletter has established itself as a federating medium to strengthen and push forward the fastly growing research community on computational developmental systems. It is thus an honor for me to be the new editor of the newsletter, and continue the excellent work of past editors, in collaboration with Juyang Weng and Zhengyou Zhang.

In this issue, you will read a profound dialog, lead by Juyang Weng and Jay McClelland, on the epistemological foundations of our field. You will also read a stimulating call for dialog proposed by Paul Fitzpatrick and entitled « Which skills most need development? ». Researchers interested in this topic are welcome to submit a response (contact paulfitz@liralab.it or pierre-yves.oudeyer@inria.fr) by August 1, 2008. The length of each response must be between 300 and 500 words (including references) due to the page limitation of the Newsletter.

-Pierre-Yves Oudeyer, Editor of the IEEE CIS AMD Newsletter

Committee news

- Prof. Jim Keller, VP Publications of CIS, presented our proposal for IEEE Transactions on Autonomous Mental Development at the IEEE Periodicals Committee meeting in Louisville, KY, on February 14, 2008. (Thank you Jim!) We have addressed most of their concerns, but some committee members are still not convinced that we will have enough papers for the new transactions. The proposal will be discussed again in June. Please send your ideas on how to convince them to AMD TC Chair Zhengyou Zhang (zhang@microsoft.com).
- AMD TC Chair has appointed Dr. Pierre-Yves Oudeyer from INRIA, France, as the new Editor of the AMD TC Newsletter. We thank the past Editor Dr. Shuqing Zeng for his excellent service to the community.
- IEEE CIS President Dr. David Fogel has approved Prof. Minoru Asada (Osaka University, Japan) and Dr. Giorgio Metta (University of Genova, Italy) as AMD TC Vice Chairs, following AMD TC Chair's nomination.
- The AMD TC Task Forces have been set up.
  Details are available at http://research.microsoft.com/~zhang/AMDTC/task.htm. Please contact the task force chair to join the task force you are interested in.
- All details of the AMD TC (organization, Newsletter, etc) are available from our TC maintained website http://research.microsoft.com/~zhang/AMDTC/
- AMD Journal Publication Repository: Early January, we sent out a call for constructing a central repository of journal publications on AMD, and we have received a large number of submissions. This is a clear testimony of the vivacity of our AMD community. The result is the repository website at http://research.microsoft.com/~zhang/AMDTC/publications-J.htm. We thank Mojtaba Solgi for his superb job in maintaining the database and creating the website.
  You can still submit your journal publication information according to the instruction at http://research.microsoft.com/~zhang/AMDTC/CallForPublications.htm

-Zhengyou Zhang, Current chair of the AMD TC

Note on ICDL Governing board election

January 2008: Four new members were elected to the ICDL Governing Board: Minoru Asada, Giorgio Metta, Brian Scassellati and Olaf Sporns. They joined the existing board members: Hideki Kozima, James (Jay) McClelland, Alex (Sandy) Pentland, Terrence J. Sejnowski, Linda B. Smith, Juyang (John) Weng (chairman) with Jochen Triesch as the communication co-chair of the ICDL Governing Board. To generate new candidates for the election, John Weng invited all past ICDL general and program chairs with terms of service. Six candidates accepted the invitation and the terms of service. They were put on the ballots with their election statements. Jochen Triesch counted the anonymous ballots from ICDL members emailed to an independent personnel. Welcome, the new members of the ICDL Governing Board!

The expanded ICDL GB elected Minoru Asada as the new chair of ICDL GB and Jochen Triesch as the co-chair.

At the steady state, the ICDL GB consists of 12 members. Every two years, the term of 4 members will expire. Each past ICDL GB member is off board for a term (2 years) before being eligible for re-election. The election of the chair and the co-chair takes place right after new members join the board.

Dialog column

How the Mind Works and How the Brain Develops


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

Jay McClelland, Department of Psychology, Stanford University Stanford, CA 94305

Studies on how the mind works can be categorized in different ways. For the purpose of this dialog, the following is a particular sequence of account:
1. *Observation of behavior*. Studying human and animals behaviors under varying task and stimulus conditions.
2. *Modeling behavior*. Computational modeling of behaviors under varying task and stimulus conditions and verification of the modeling from studies in category 1. This level of modeling does not necessarily take into account how the mind develops from experiences or how the brain works.
3. *Modeling the brain basis of behavior*. Computational modeling of how the brain works -- i.e., how its activity gives rise to behavior, at many different scales, including the cortex, specific circuits and individual neurons. This level of modeling does not necessarily take into account how the brain develops from experiences.
4. *Modeling brain development*. Computational modeling of how the brain develops from conception to adulthood, across different scales. This level of modeling takes into account how the brain works and how the brain develops, and considers how experience and context structure brain development and produce the functional consequences of this development.

Three major dimensions are involved in the above order: from observation to modeling; from the mind to the brain; and from brain to the experience-dependent developmental processes that shape both brain and behavior. Of course, one could proceed through these dimensions in a different order; for example, one could develop models of how experience shapes behavior, without regard the brain basis of this process; and surely, experimental investigation of the brain (and of its development) must complement computational studies at levels three and four.

The questions that we would like to raise for discussion are these:
a. First, can a particular level of investigation benefit from the other levels of investigation?
b. Second, do we see a light in the tunnel? In other words, does computational modeling at level 4 potentially unify major studies across all four categories? For example, will Bayesian modeling at a behavioral level be explained or modified by studies in neuronal computation and development?c. Are cell-based computations within the networks in the brain too basic to give rise to rich accounts about how the mind works?
d. Can there be a computational theory (in the sense of Marr [1]) that encompasses investigations at all four levels? Marr himself considered perception and its neural basis but did not give much attention to experience or development.

REFERENCES
[1] D. Marr, Vision, 1982, San Francisco: Freeman.

The minds start to work when the brains develops

Minoru Asada Dept. of Adaptive Machine Systems, Graduate School of Engineering, Osaka University, Suita, Osaka 565-0871, Japan

The approach we are taking in cognitive developmental robotics [1] (hereafter, CDR) seems similar to the one in the dialog by Juyang and Jay, but a little bit different. Roughly speaking, CDR consists of three areas: developmental psychology, computational neuroscience, and intelligent robotics; the mutual feedback among these areas is essential in CDR. Therefore, ideally, we should start the study considering the relationship among these areas from the beginning. That is, when observing behaviors, we consider the computational model and its realization. Rather, based on these kinds of models, the observation could be meaningful. Of course, that depends on how complete they are. If the model has been almost established, the observation can be the verification phase of the model. On the other hand, if the model is far from complete, just at the level of a simple idea, the role of the observation is to shape the idea. In any case, the observation should always have the model of any level in CDR, and therefore the order from 1 to 4 in the dialog does not seem meaningful. The answer to the question "can a particular level of investigation benefit from the other levels of investigation?" generally depends on the level, but it should do because CDR aims at interactions in any level to reveal the mystery and approach to it. We should not expect that the existing models or theories work very well for our purpose.

We should build our own theory or model. Of course, the existing ones can be good tools and/or starting points to build our own. Bayesian modeling is a typical example of the excellent existing tools, but it is not perfect to explain and to reproduce the developmental process we are focusing on. Rather than general model construction, we should start one case study with deep insight, and try to extend it to other cases.

The scale problem between a microscopic level such as cell-based computations and a macroscopic one such as behavior or mind level is one of the fundamental issues. The model we'd like to construct is expected to solve the issue, but again, we should start a case study first, then try to extend it.

Stumbling Towards Integration

Gedeon O. Deák, Dept. of Cognitive Science,UCSD

The history of science repeatedly shows that single disciplines hit roadblocks in the search for answers to difficult questions, and interdisciplinary traffic is sometimes the path to new theoretical and empirical insights (think Newton or Darwin; materials sciences or neurosciences). Weng and McClelland's questions tap into a larger metascientific question: How do individual scientists identify viable "targets of opportunity" for interdisciplinary effort, and bring to bear the right methodologies to make progress?

The case in point is how to integrate across levels to explain the development of physically (encompassing neurally) embodied complex behavior. The practical re-statement of the meta-question is: "What interdisciplinary opportunities should be exploited *now*?" Many possibilities are intractable, unethical, or impossible. Excluding these, let us assume that interdisciplinary efforts are intended to circumvent the common difficulties of each discipline. What are those limitations, for the intersection of human behavioral and cognitive development, systems neuroscience, and dynamic modeling/machine learning?

With respect to modeling embodied agents some familiar problems are: scaling up toy models to biologically plausible levels of complexity; avoiding hand-wiring and over-fitting; ensuring neural plausibility of machine learning techniques. A major difficulty is the *dual modeling problem* (Deák, Bartlett & Jebara, 2008); that is, believable simulations require not only a neurally inspired, biologically feasible learning agent, but also a plausible representation of the learning environment and input to the agent. With respect to systems neuroscience, much evidence comes from animal models, but precise human data on many topics is limited. This is exacerbated for developmental data: FMRI methods, for example, are quite difficult to use with young children; electrophysiology is also limited due to testing demands.

Important advances are occurring in developmental structural MRI and DTI projects, and these will inform a wide range of developmental theories. However, neuroscience has its own scaling problems: moving from molecular to cellular to systems-level descriptions, we lose precision in describing complex representations, and the dynamics by which they emerge in infants and children. Of course, adding epigenetic or emergent dynamics to a model will, all else being equal, increase the complexity and parameters to be simulated. The problem lies in deciding which simplifications of the dynamics are least problematic at the outset.

With respect to cognitive/behavioral experiments, developmental studies traditionally collect sparse data, with limited information about, for example, variance in response speed, the nature of errors or non-responses, physiological indices (e.g., sympathetic nervous system activation symptoms), or perceptual-motor dynamics. This is changing as more labs integrate, e.g., heart rate data with looking-time studies, or eye movements with studies of word or sentence processing. A worse difficulty is that some experiments test theories that are not well specified, or not informed by neuroscience. Certainly there are many practical difficulties in recruiting, testing, and coding data from infants and young children.

With these problems in mind, what can computational models bring to the study of human development? One is proof of concept: for example, Lewis and Elman [3] used simulations of recurrent neural networks, given as input sentences from a child-language corpora database, to show that one of Chomsky's "poverty of stimulus" cases (i.e., syntactic structures argued to be unlearnable) was in fact learnable. Because Chomsky claimed unlearnability by fiat (i.e., intuition), evidence from even a simple ANN would trump his claims. The point is that even a "toy" simulation, applied to the right theoretical questions, can be useful.

Computational models also are excellent for debunking "magical" thinking about infant abilities. Developmentalists have sometimes attributed infant perceptual sensitivities to "innate knowledge," although the meaning of "innate" is not often specified. So the claim that infants innately prefer and process human faces must be explained; simulations such as by Dailey and Cottrell [2] suggest that fairly general constraints on neonatal visual processing could cause units in a neural network to develop face-recognition specialization, given the visual information in faces. More recently, Butko, Fasel and Movellan [1] showed that a minimal amount of natural visual and auditory input would yield face-finding and classification by a robot "baby."

Because the constraints are well-specified in computational implementation, they often yield predictions to be falsified by behavioral studies. For example, Triesch, Teuscher, Deák and Carlson's [4] reinforcement-learning simulations suggest that speed of habituation will modulate the acquisition of gaze-following by infants; this prediction is being tested in a longitudinal behavioral study.

One point is that dual-modeling problems hit behavioral research as well as modelling. That is, if a theory calls some capacity "innate" (i.e., invariant phenotype), falsification entails a description of the input to the organism from some early point in development of the functional system.

What about the intersection of neuroscience and developmental theory? One issue is that the neurosciences have expanded so rapidly in recent decades that it is challenging for a behavioral researcher to learn which advances might be relevant--for example, what research in genetic co-modulation of pyramidal dendrite morphology should constrain one's models of early spatial representations? Another challenge is that it is hard to design ethical tests of experiments to pit one neurodevelopmental account against another. Although advances in high-density EEG and MEG hold promise, it is important to recognize that infants and children are sort of "hyper-embodied thinkers"—moreso than adults, their cognitive processes are externally reflected in their motor activity (visual, fine motor, gross motor) and are highly influenced by state and affect. Moreover, their responses are highly dependent on testing materials. Thus, the kinds of methodological shortcuts we use for experimental control—and convenience—when testing adults (e.g., showing repetitive static images on a monitor, while restricting head motion) might dramatically affect young children's test performance. Thus methods for developmental cognitive neuroscience, using any technology (EEG/ERP, fMRI, MEG, NIRS), cannot just be simplified from adult studies. They must be re-conceptualized and re-engineered for different age groups.

All of these concerns push us towards a conclusion: for interdisciplinary research to have impact and drive theoretical progress, true interdisciplinary collaborations are needed. Yet such collaborations require not only common purpose (i.e., an exciting, compelling problem) but a cultural of willingness to learn another discipline, and eventually a body of shared knowledge that facilitates conversation. Such interactions may require renovating academic and funding cultures.

Universities sometimes rigidly reinforce paradigmatic disciplinary research, and even impose economic and professional barriers to interdisciplinary work. Yet just as professionals in many disciplines are required (and paid) to obtain continuing education, universities can give faculty teaching releases to take courses to attain expertise in a second discipline. Also, funding agencies can start to "see the light"—for example, NSF's Human Social Dynamics competition required every proposal to include PIs from at least two disciplines.

To re-frame Wang and McClelland's question more pointedly, can developmental science survive if interdisciplinary efforts do not become the norm? Finding the right problem and the right collaborative "blend" are unavoidable hurdles, but those processes will be facilitated if the leaders in each field make efforts to convince universities, publishers and funding agencies to support these efforts.

REFERENCES
[1] N.J. Butko, I.R. Fasel, & J.R. Movellan. Learning about humans during the first 6 minutes of life. Proceedings of the International Conference on Development and Learning, 2006.
[2] M.N. Dailey, G.W. Cottrell. Organization of face and object recognition in modular neural network models Neural Networks 12, 1999, 1053 – 107.
[3] J.D. Lewis, and J.L. Elman, A connectionist investigation of linguistic arguments from poverty of the stimulus: Learning the unlearnable, Proceedings of the 23rd Annual Conference of the Cognitive Science Society, 2001, 552-557.
[4] J. Triesch, C. Teuscher, G. Deák, and E. Carlson, Gaze-following: Why (not) learn it? Developmental Science. 9, 2006, 125-147.

Vision is Hard Enough for Me!

John K. Tsotsos, Dept. of Computer Science & Engineering and Centre for Vision Research
York University Toronto, Ontario, Canada

The authors' Dialog Initiation proposes a series of steps or tasks towards the goal of understanding the development and function of the human brain and mind. They raise very important issues; but the scope of the work is far beyond what I feel comfortable commenting on. So I will focus where I might have a bit of competence. In order to understand behavior, observation of that behavior is a sensible thing to do. What is not as clear is what we wish to observe nor how we wish to observe it. There seem to be at least two strategies, one far more common than the other. The bulk of current experimental work focuses on detailing some aspect of human performance by varying experimental parameters or by expanding experimental scope. Something interesting is observed and then its characteristics are explored in detail. Experimental methods to maximize the value of the exercise have been detailed and proven over time [7]. The less common method has been to develop a theory with predictive power; testing falsifiable predictions then leads to theory refinement and further predictions. I certainly see a role for the former strategy but advocate a greater use of the latter, in conjunction with the former. It is also important to remember that not all observable behavior may lead to conclusions about function; several authors have argued this previously (see [10, 3] for visual search, [1] for more general sensory experiments).

Modeling also takes many forms. One can try to fit curves to data for example and thus there is predictive power by interpolation or extrapolation (Bundesen's theory of visual attention is an excellent example [2]). But this approach does not lead to a theory that can generate the behavior (as in a robotic camera system, for example). One can view the brain as a black box and only care about its input-output behavior. Learning methods can then find statistical correlations between input and output, creating feature vectors that a classifier can then cluster; but this may tell us nothing about what's inside the box, i.e., the brain. However, these correlations can be part of a functional system that can reproduce the behavior [8]. In an engineering sense this has value, but as a science of the brain I am not so certain. Finally, one may attempt to develop a first principles theory, such as my own for visual attention [10], where the 'first principle' is search optimization and basic proofs come from computational complexity. This theory leads to a host of counterintuitive predictions many with strong support now, as well as to systems that demonstrate the desired behavior [13, 15]. First principles theories are necessary and play foundational roles in other sciences: why not in brain science?

The above paragraphs are motivated by my own experience, in vision, in AI and some robotics. But before we give Marr all the marbles on computational theory and perception at a neural basis we should note that he explicitly excluded any aspect of perception that required more than 160ms of response time, i.e., anything that requires scrutiny, attention, reasoning and more (see p. 96, [6]) - anything that is difficult!

His views on computational theory had value because he took what was already part of basic computer science (which presumably he was just introduced to) - top-down software development - and eloquently showed the human vision community that those same ideas had value in modeling. The ideas were not really new to those of us who came from computer science to vision. Moreover, and specifically because he was not a computer scientist, he neglected a critical component - tractability - from his perspective (see [10]). Brooks neglected this too when making his claims about his version of behaviorist robotics extending to human behavior [14].

In both cases, this is a fatal flaw. It is relatively easy to propose a solution that may suffice 'in principle'; but it is a different matter altogether to ensure that it is realizable in some implementation (brain or machine). Ensuring low computational complexity of the problem being solved is one step (the brain is not solving intractable problems); and, the computational complexity of the algorithm also must be tractable before an implementation is considered.

What about computational models of development that must obviously include vision? Consider the following claim of the AMD paradigm [16]: "For example, the visual world seen by our eyes is very complex. The light that falls on a particular pixel in a camera depends on many factors--lighting, object shape, object surface reflectance, viewing geometry, camera type, and so on. The developmental principles capture major statistical characteristics from visual signals (e.g., the mean and major directions of signal distribution), rather than every aspect of the world that gives rise to these signals.

A task-specific programmer, in contrast, must study aspects of the world around the specific task to be learned; this becomes intractable if such a task, such as vision, speech, or language, requires too many diverse capabilities."

Why highlight this? Because it is exactly the opposite of what we see theoretically! That is, without task-specific knowledge, vision (and I believe any sensory process) has potentially exponential complexity while the addition of even small amounts of task information makes it linear [11, 9] in both worst and median case analyses. It seems to be a cornerstone of AMD program that task is open, that robots are not task-specific. But care should be taken so that although the robots should be generic and thus not single-task devices, they need to be able to use task information effectively. I worry that the methods used truly eliminate this ability by burying any knowledge in a mass of statistics.

Let's make sure we are all using the term 'tractable' to mean the same thing; I am using it in the sense that tractability is equated with polynomial computational complexity and intractable means exponential complexity. And there are indeed polynomial problems where the variables have large values and thus they seem inefficient just like there are exponential problems where the variables take on small values and thus they may seem efficient. It's not so straightforward in other words. The task of statistical learning I am guessing has exponential complexity because so many learning algorithms in the past have also been shown to have exponential behavior [4,5]. Saying that the new AMD algorithms are tractable and will ignore just the right bits and keep only the important bits seems to me to assume too much; we don't know what those "just right" representations are!

This is not to say that the brain is solving any sort of intractable problem; rather it does say that the brain is not solving the generic problem of vision, that it has re-shaped that problem to turn it into a tractable one, and that one of the elements that it uses is task knowledge when it's available. We know more about the visual system than any other part of the brain. Yet, our theoretical understanding, our models of all kinds, is very far from satisfactory. There is an enormous amount still to be done both theoretically and experimentally. So the question is: will the procedure outlined by Weng & McClelland lead to a deep enough understanding of mind (and thus vision as well) so that it can contribute to that understanding or will it at best yield a superficial sensory ability for these simple robots? If the former then I am excited by the prospect. If the latter, then I'll stick to my current strategy. Vision on its own is hard enough!

REFERENCES

[1] G.S. Brindley. G.S. Physiology of the Retina and Visual Pathway, (chapter 5, pages 132–138), 1990, Williams & Wilkins.
[2] C. Bundesen, A theory of visual attention. *Psychological Review*, 97, 523-547, 1990.
[3] L. Chelazzi, Serial attention mechanisms in visual search: A critical look at the evidence, *Psychological Research* 62: 195-219, 1999.
[4] S. Judd, On the complexity of loading shallow neural networks, *Journal of Complexity*, Volume 4 , Issue 3 (September 1988) p: 177 – 192, 1988.
[5] S. Judd, Neural Network Design and the Complexity of Learning, 1990, MIT Press.
[6] D. Marr, Vision: A Computational Investigation into the Human Representation and Processing of Visual Information, 1982, Henry Holt and Co., Inc. New York, NY, USA.
[7] N.A. Macmillan, C.D. Creelman, Detection Theory: A User's Guide, 2005, Routledge.
[8] J. Mutch, D. Lowe, Multiclass Object Recognition with Sparse, Localized Features, IEEE Computer Vision and Pattern Recognition

[9] R.A. Rensink, A new proof of the np-completeness of visual match. Technical report, Computer Science Department, University of British Columbia, 1989.
[10] J.T. Townsend, Serial vs. parallel processing: Sometimes they look like tweedledum and tweedledee but they can (and should) be distinguished. *Psychological Sciences*, 1, 46-54, 1990.
[11] J.K. Tsotsos, The complexity of perceptual search tasks. In Proceedings of International Joint Conference on Artificial Intelligence, pages 1571–1577, 1989.
[12] J.K. Tsotsos, Analyzing Vision at the Complexity Level, *Behavioral and Brain Sciences* 13-3, p423 – 445, 1990.
[13] J.K. Tsotsos, S. Culhane, W. Wai, Y. Lai, N. Davis, F. Nuflo, Modeling visual attention via selective tuning, *Artificial Intelligence* 78(1-2), p 507 – 547, 1995.
[14] J.K. Tsotsos, On Behaviorist Intelligence and the Scaling Problem, *Artificial Intelligence* 75, pp. 135 – 160, 1995.
[15] J.K. Tsotsos, Y. Liu, J. Martinez-Trujillo, M. Pomplun, E. Simine,K. Zhou,, Attending to Visual Motion, *Computer Vision and Image Understanding*, Vol 100, 1-2, p 3 – 40, 2005.
[16] J. Weng, J. McClelland, A. Pentland, O. Sporns, I. Stockman, M. Sur, E. Thelen, Autonomous Mental Development by Robots and Animals, *Science*, Volume 291, Number 5504, Issue of 26 Jan 2001, pp. 599-600, 2001.

A Comprehensive Approach to Studying How the Mind Works and the Brain Develops

Jeffrey L. Krichmar, Department of Cognitive Sciences, University of California, Irvine, CA 92697

Studying how cognition, perception, and consciousness emerge from neural mechanisms is one of the open frontiers of science. Although much progress has been made over recent years, there are many unanswered questions on how the brain gives rise to the complex behavior we associate with biological organisms. I applaud Juyang Weng and Jay McClelland for initiating this discussion. Success will require research that reaches across disparate fields such as cognitive science, medicine, neuroscience, psychology, mathematics, engineering, and computer science. A promising future initiative, the Decade of the Mind, proposes a cross-disciplina ry approach to better heal, understand, enrich, and model the Mind [1].
In my opinion, understanding how the mind works and how the brain develops will require a multifaceted approach that extends beyond observation and modeling. It will require progress in behavioral neuroscience, electrophysiological recordings, brain imaging, and psychophysics. Comprehension of the mind's inner workings will require new tools that can deeply probe mental processes. Research is needed to develop techniques that can record simultaneously from very large numbers of neurons, over multiple brain areas, in behaving animals. Current brain imaging in humans (i.e., fMRI, PET, EEG, and MEG imaging techniques) do not have sufficient spatial or temporal resolution to completely examine the dynamics of the brain in action.

Because of the scope of the problem, theoretical and computational modeling at multiple levels will be crucial for understanding how the mind works. Models have explanatory and predictive power. For example, a recent model from our group predicted that the perforant path from the entorhinal cortex to the CA1 subfield of the hippocampus is important in the recall of episodic memories [2, 3]. However, this prediction, and any prediction from a computational model, needs to be confirmed through experiments with animal models. Moreover, models must be constructed from empirical data. A major challenge of the future will be synthesizing the various levels of computational modeling with empirical data, and consolidating this information into a coherent picture of the mind. The construction of future computational models of the nervous system will require more accurate and comprehensive empirical data. Databases are becoming available to meet this need [4]. However, a complete functional description of the pathways and connections in the nervous system will be necessary for constructing accurate models [5].

Recent progress in cognitive science, computer science, engineering, medicine, neuroscience, and psychology make truly understanding "How the Mind Works and the Brain Develops" a real and exciting possibility. Computation theory and modeling is just one component, but a very important one, which will be necessary to ensure success in this endeavor.

REFERENCES

[1] Albus, J. S., Bekey, G. A., Holland, J. H., Kanwisher, N. G., Krichmar, J. L., Mishkin, M., Modha, D. S., Raichle, M. E., Shepherd, G.M., & Tononi, G., 2007, Science (New York, N.Y 317, 1321.
[2] Fleischer, J. G., Gally, J. A., Edelman, G. M., & Krichmar, J. L., 2007, PNAS 104, 3556-3561.
[3] Krichmar, J. L., Seth, A. K., Nitz, D. A., Fleischer, J. G., & Edelman, G. M., 2005, Neuroinformatics 3, 197-222.
[4] http://ndg.sfn.org
[5] Sporns, O., Tononi, G., & Kotter, R., 2005, PLoS Comput Biol 1, e42.

Response to « **How the Mind Works and How the Brain Develops** »

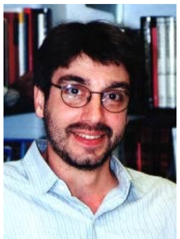

Olaf Sporns, Department of Psychology, Indiana University, Bloomington, IN 47405

John and Jay are asking a series of interrelated questions about brain, mind and development that are of interest to all of us in the AMD and ICDL communities. Their central question concerns the linkages between different levels of investigation in the study of brain, cognition and behavior and how these linkages may be useful for establishing a unified theoretical framework for brain and mind.

Personally, I am quite optimistic. We have seen great progress in biological approaches and in computational/modeling approached in recent years. I am encouraged by an increasing realization on the side of empirical neuroscientists that the brain must be regarded as a dynamically coupled system, as a system even that is embedded in its environment. In other words, the purely reductionist research strategy that eliminates mental and developmental process entirely in favor of exclusively single-cell or molecular neural correlates has not delivered a complete understanding of the brain, let alone the mind, and an alternative more synthetic and systems-based approach to the brain is on the upswing. On the computational and theoretical side there are numerous new developments as well that link biological processes to computational principles, for example in neural coding, plasticity, brain dynamics, functional neuroimaging and embodiment. A challenge has been to unify these approaches in joint research projects and to make the joint advances public and visible. Here is where the AMD/ICDL community fulfills a significant role for the continued progress in combined biological/computational approaches to development.

John and Jay ask if "cell-based computations within the networks in the brain [are] too basic to give rise to rich accounts about how the mind works"? I understand the spirit of the question, and I do not think that mind can be reduced to the functioning of single neurons, even if we could observe and record all of them. Clearly there are emergent and large-scale properties that become manifest in cognition but are entirely absent when one considers single neurons. A friend of mine once said that trying to understand the mind by analyzing the activity of a single neuron is a bit like attempting to understand the US economy by following around a single shopper in a supermarket. That is not to say that observations of shoppers in supermarkets are not useful! But it cannot be all we're doing. We need to think hard about how we can bridge the levels that John and Jay are illustrating. I think there is tremendous promise in the study of development, because development is clearly of central importance to both brain and mind and a better understanding of the computational principles of development may get us a lot closer to finding common ground.

Do we see a light at the end of the tunnel? I think yes (and I don't think it is a train coming our way!)

Reverse Engineering the Brain*

Paul Werbos, National Science Foundation, Suite 675, 4201 Wilson Blvd.
Arlington, VA 22230.

On February 19, 2008, the National Academy of Engineering listed fifteen grand challenges for the coming century, including: "Reverse-engineering the brain, to determine how it performs it magic, should offer the dual benefits of helping treat diseases while providing clues for new approaches to computerized artificial intelligence."

People have talked about understanding how intelligence emerges in the brain, and building artificial intelligence, for decades and decades. For example, in the 1960's, one of the two leading centers on artificial intelligence promised NASA that they would deliver a full-fledged artificial human brain, based on linguistics and symbolic reasoning, in a robot ready to be sent to Mars in the 1980's. Likewise, in the past, many neuroscientists suggested that we could reach an understanding of intelligence and learning in the brain by first doing exhaustive probes of simple organisms like snails (aplysia), and extracting the mathematical principles; however, current research on such organisms suggests that much of their behavior is indeed genetically programmed, and is not so helpful in understanding how we can learn new things above and beyond what is already coded into our genes.

Yet another stream of research has stressed that the human mind has an almost magical ability to understand things beyond the scope of conventional computing – and then suggested that we follow up by building "computers to compute the noncomputable."

Engineering today has made enough progress to allow us to construct a fourth pathway of research, which can allow us to bypass the obstacles which have limited the accomplishments of the old paths. The old paths still have much to contribute, but something new is needed, to achieve a real functional understanding of intelligence in the brain. In order to meet this new opportunity, we will need to focus on the unique new directions in engineering which make it possible, and to focus on ways to exploit these new directions. It is necessary to build strong partnerships with empirical neuroscience (especially systems neuroscience) and psychology, without falling into the limitations of traditional research in those fields.

The first key goal on this new path is to focus on trying to understand learning-based general intelligence, as it exists in brains ranging from the crudest fish to the smallest mammal. While others focus on snails or try to build an "artificial Einstein," the new opportunity lies at this middle level. Symbolic reasoning and empathy ("mirror neurons") have given great power to the human brain, not shared by these lower organisms. Yet before we can really understand the "deep structure" which underlies language and semantics, we must first understand the functional wiring and learning which human brains and mouse brains share. The six-layer neocortex of the human brain is almost indistinguishable from that of the mouse. Different topographic regions in the human brain usually contain different knowledge from what mouse brains contain, but human brains differ from each other, and can change with time; the famous "edge detectors" hard-wired into certain regions are easily relearned in other parts of the brain, when the system is damaged, if the wiring from thalamus to neocortex is not cut.

The next key element on this new path is to focus on cognitive optimization and prediction. General intelligence is not the ability to do well in a narrow specific task, like playing a good game of chess. Computer scientists have shown how that kind of narrow focus is a dead-end, when it comes to learning about intelligence in general. On other hand, the task of adaptive, anticipatory optimization in the face of a large, unknown, nonlinear stochastic environment is of general importance. It is a task where engineers have led the development of new types of mathematical design, which can handle far more complexity than traditional forms of reinforcement learning which have been popularized in computer science. The most important research questions here are: (1) how can we deepen our understanding of the mathematical principles of adaptive, anticipatory optimization and of prediction, so as to cope with the kind of complexity in time and in space which the brain can handle?; (2) how can we insert more brain-like approaches into engineering technology?; (3) how can we unravel the circuits and methods which the brain uses to learn to predict or to optimize better, in a way which transcends specific testbeds for prediction or optimization? How can the smallest mouse learn to maximize its probability of survival?

As an example, Nicolelis and Chapin showed in the 1990's that certain cells in the thalamus predict the sensor inputs that are received by other cells, and relearn to predict well after damage to the system. How does this finding apply across all sensory inputs, and how does the brain learn this?

In general, the goal here is not to **predict** brains in a "behaviorist" or "positivist" way, as if they were like a passive crystal of ice; rather it is to understand the **general functional capabilities** of brains, so that we can better understand, work with and enhance the intelligence possessed even by the smallest mouse. How do **brains** themselves learn to predict better and better with time? Of course, one cannot do full justice to the many vicissitudes of these topics in such a brief statement. My apologies to many others who deserve far more attention. For broader and more elaborate discussions of various aspects of this, google on "COPN neural", look at www.werbos.com, or ask me to send you the paragraph on "AIS" (Adaptive and Intelligent systems) which is looking for more new proposals related to the mathematical and engineering issues here.

Reply and Summary: How the Minds Works Results from How the Brain Develops

Juyang Weng and Jay McClelland

We have received an array of very interesting responses, covering a very wide spectrum.

To facilitate the following discussion, here are the four levels that we outlined in our dialog initiation [1]:

- Level 1: *Observation of behavior*. Not necessarily including computational modeling.
- Level 2: *Modeling behavior*. Not necessarily including brain basis, but, in a developmental context, may consider the role of the environment.
- Level 3: *Modeling the brain basis of behavior*. Not necessarily including development.

- Level 4: *Modeling brain development*. This level underlies all the others and, we argued, requires incorporating a role for experience and environment.

Asada agreed with our argument for the inclusion of knowledge of developmental psychology and computational neuroscience in cognitive developmental robotics. However, he advocates "we should start one case study with deep insight, and try to extend it to other cases." Even though a case-by-case modeling involves changes from infancy to childhood to adulthood, the does not necessarily seek brain basis. Much work in developmental robotics currently is at this level 2. This level is certainly useful in many ways. Our argument is that ultimately full understanding will depend on going beyond this level. In reply, a typical argument that we have heard many times is something like: "well, robots do not have to duplicate the brain, and anyway, not much is known about the brain and even less about its development." However, we argue that without understanding how the brain develops (knowledge at level 4), we will not be able to truly understand how the mind works, and developmental robotics will not be able to realize its true potential. We ask: what can our community do to facilitate interdisciplinary understanding and collaboration so that the insights that will come from incorporating level 4 have the best chance of facilitating progress in developmental roboticists?

Deak gives an illuminating account of the benefits and challenges of inter-disciplinary research. His "innateness" examples are especially illustrative for the benefit of interdisciplinary research, between developmental robotics and developmental psychology. However, he raises a concern about research that might be seen as claiming to address the brain basis of development: "The current network models have not provided convincing links with the brain basis. Thus, if a network learns a behavior, we still do not know whether the brain learns the behavior." Certainly, it is not enough just to say that one's network model attempts to address the brain basis of development. Stronger constraints from developmental neuroscience are necessary, and, in our view, welcome. To truly understand the causality about how the mind works, the research community needs studies about the brain basis (level 3) and brain development (level 4) to be more fully integrated into the discipline of developmental robotics.

Tsotsos studied the important issue of complexity of visual search. If we are correct, Tsotsos' "first principle" targets at level 3 in the hierarchy that we outlined. If level 4 knowledge is taken into account (e.g., neuronal migration, branching, connection and adaptation in cortical development are all environmental dependent), Tsotsos might like to rethink about his attention models, and so might Olshansen, Anderson and Van Essen1993 [2]. At level 3, the brain might appear as a static, very capable machine whose physical complexity is beyond what could be achieved by hand design and whose computational complexity is beyond comprehension; but at level 4 the brain is a gene-regulated biological machine that self-organizes into an amazingly capable and efficient machine through developmental experiences. Any model that requires cells to have a totally gene-specified connection patterns is likely to be wrong at level 4, as mounting evidence from developmental neuroscience has demonstrated (e.g., traveling waves of activity in the prenatal retina help organize the development of neuronal connections to the thalamus). Does this mean that at level 4 the brain does not use task-specific information as Tsotsos questioned? Of course not. But in our view, most task-specific information is not in the genes; it is acquired from environment. Tsotsos might be very interested in looking into how visual attention skills are learned while the brain wires and constantly rewires itself. Of course, we do not have complete solutions at level 4. Welcome to this community, John!

The responses from Krichmar and Sporns are quite optimistic and point to the importance of modeling development. Sporns mentioned a great analogy about observations of shoppers in supermarkets as opposed to understanding the US economy. Sporns and Krichmar, here is a hard ball from Tsotsos: "We don't know what those `just right' representations are!"
As we can see, two approaches exist to taking the ball from Tsotsos: (a) Level 3 approach: The representation is considered static in the scope of a level 3 study). (b) Level 4 approach: the representation is constantly and dynamically changing, by the cell-centered developmental mechanisms (manifested by interactions of many cells) and the cell's environment (i.e., the brain's internal and external environments). As Sporns implied, without insightful investigations at level 4, even if there is a light in the tunnel, we cannot get to it.

Werbos' response provides an excellent account of the knowledge and lessons learned in the past. Linguistics and symbolic reasoning in Werbos' story roughly belongs to level 2. Werbos' *aplysia* story tells us that genetically programmed behaviors are not so helpful to our understanding of how the human brain learns new things. However, we may imagine that even if the object of study is changed from aplysia to *homo sapiens*, the "passive crystal of ice" view of the brain, as Werbos called it, is not going to lead us a practical capability of "reverse engineering the brain."

This points to the limitations of level 3 studies, although studies in all levels are useful. Werbos emphasized both a brain basis and a learning basis for intelligence, which what our level 4 is all about.
We propose the following two observations: One on *the brain causality hierarchy* and the other about "*light in the tunnel*".

a. Brain causality hierarchy: All the four levels of studies have much to contribute, but the higher the level, the deeper understanding about the causality of the lower order levels and more systematic in engineering of machine intelligence. For example, level 4 explains the mechanisms of brain development that give rise to the brain basis of behaviors that are captured in static computational models (levels 2 and 3). A learned adult brain is extremely complex. A case-by-case study of learned brain basis of behaviors (level 3) is a study of the result of development, but a corresponding study at level 4 is a study of the causal processes that lead to such a result. The mechanisms of causality at level 4 also systematically apply to investigations at level 2.
b. "Light in the tunnel": The current popular Bayesian approach to understanding behavior is usually applied at the lower-numbered levels (1 and 2 primarily). Bayesian techniques have been fruitfully used in engineering and mathematics. They can also be used to model human children's classification and other behaviors, and some Bayesian researchers have begun to use these approaches to understand changes in children's abilities based on experience. However, Bayesian models have not been mapped to the brain basis --- neuronal computation using physiologically supported cell computation mechanisms. Neither does it model how the brain development using physiologically supported cell mechanisms for migration, connection, and learning. On the other hand, learning models using physiologically supported cellular learning mechanisms (e.g. Hebbian learning or its more detailed version Spike Time Dependent Plasticity (STDP)) is at level 4, but such work often fails to address the emergent phenomena seen at higher levels. There is exciting ongoing work at both of these levels, but little is known about the relationships between them. We argue for bringing these different approaches together to see if we can find our way to the light in the tunnel.

We would like to thank all the authors who have contributed their responses. Regardless of the positions in authors' views, their contributions greatly enriched the discussion about this important subject.

REFERENCES

[1] J. Weng and J. McClelland, "How the Mind Works and How the Brain Develops," *IEEE CIS AMD Newsletter*, vol. 4, no. 2, 2007. http://www.cse.msu.edu/amdtc/amdnl/

[2] B. A. Olshansen, C. H. Anderson and D. C. Van Essen, "A neurobiological model of visual attention and invariant pattern recognition based on dynamic routing of information," *Journal of Neuroscience*, vol. 13, 11, pp. 4700-4719,1993.

Dialog Initiation

Which skills most need development?

Paul Fitzpatrick, RobotCub humanoid project, University of Genoa, Italy.

We know that as adults, every skill we possess arose through an intricate developmental process of interlocking behaviors, innate and learned. Robot abilities are not generally constructed this way, although our community is doing its bit to change this. Are there skills for which the case for development in implementation (for robots) or description (for humans) is particularly strong? Without development, what skills can reasonably be implemented or described, and which skills will lie outside our reach?
As a roboticist, I speculate that the skills most in need of developmental implementation are those that retain a strong need for plasticity in their mature state, since a cheap and powerful way to implement such plasticity would be to retain parts of the developmental process used to create the skill in the first place. On the other hand, skills which don't retain this kind of plasticity in their mature form are candidates for abstraction and implementation without appeal to development.
Consider obstacle avoidance, one of the most frequently implemented abilities for any new robot, where the robot sidles its way around objects in the environment. To a first approximation, new objects or arrangements of objects don't call for much plasticity, since they can probably be evaded in much the same way as other objects. And in fact this ability has been implemented satisfactorily without developmental methods (although it is never quite as robust as one might like).
In contrast, consider grasping and manipulation. These "hands-on" skills depend heavily on the details of the objects being manipulated, and have so far been satisfactorily implemented only in scenarios in industry, entertainment, etc. where the

robot interacts with a theatrical version of the world, with everything laid out just so by a team of human servants.

I suspect that grasping and manipulation are skills whereby an attempt to implement the mature skill with sufficient robustness in unstructured environments, through fall-back grasping strategies and active probing, would be of the same order of complexity as a developmental implementation.

My questions: what skills do people see as most in need of a developmental implementation (for robots) or description (for humans), and why? Are there skills which are "killer applications" for developmental robotics, that we just can't imagine implementing in any other way?

Responses should be sent to paulfitz@liralab.it or pierre-yves.oudeyer@inria.fr by August 1st 2008.

Calls for participation

Eighth International Conference on Epigenetic Robotics: Modeling Cognitive Development in Robotic Systems

The 8th edition of the International Conference on Epigenetic Robotics (EpiRob) will be held at the University of Sussex in Brighton (UK) on July 31-Aug 2, 2008.

In the past 7 years, the Epigenetic Robotics annual conference has established itself as a unique place where original interdisciplinary research from developmental sciences, neuroscience, biology, cognitive robotics, and artificial intelligence is being presented. Psychological theory and empirical evidence is being used to inform epigenetic robotic models, and these models can be used as theoretical tools to make experimental predictions in developmental psychology.

As a special feature, Epirob'08 highlights a specific organizational theme: evolution and development as related processes of change. EpiRob'08 is held just before ALifeXI (August 5-8) in nearby Winchester and we hope this proximity in space and time will promote cross-fertilisation between both communities. EpiRob'08 will feature three distinguished invited speakers: Eva Jablonka, Susan Oyama and Domenico Parisi. This list will be completed with a local developmental biologist to be announced.

Web : http://www.epigenetic-robotics.org

7th International Conference on Development and Learning

The 7th International Conference on Development and Learning will be held at Monterey, California, on August 9th-12th 2008.

The scope of development and learning covered by this conference includes perceptual, cognitive, behavioral, emotional and all other mental capabilities that are exhibited by humans, higher animals, artificial systems and robots. Investigations of the biological and computational mechanisms of mental development are expected to improve our understanding of the working of the whole range of mental capabilities in humans and to enable autonomous development of these highly complex capabilities by robots and other artificial systems. The International Conference on Development and Learning strives to bring together researchers in neuroscience, psychology, artificial intelligence and robotics and other related areas to encourage understanding and cross-fertilization of latest ideas and results from the different disciplines.

This year, ICDL will feature the following three distinguished speakers : Richard Aslin, Terry Jernigan, Andrew Ng.

Web : http://icdl08.org

Editor : Pierre-Yves Oudeyer Editorial Assistants : Adrien Baranès and Zhengping Ji