AI Magazine Volume 23 Number 2 (2002) (© AAAI)

Autonomous Mental Development

Workshop on Development and Learning (WDL)

Juyang Weng and Ida Stockman

This article describes a workshop on mental development and learning issues that are relevant to both machine and human sciences. It was jointly funded by the National Science Foundation and the Defense Advanced Research Projects Agency and held at Michigan State University on 5 to 7 April 2000.

The workshop was motivated by the budding research area that is dedicated to the idea of building *autonomous developmental robots,* that is, machines that develop mental skills through online, real-time interactive experiences in the real physical environment. At the same time, developments in the human sciences are opening up new possibilities for studying humans in ways that can benefit the goal of creating more intelligent machines. New findings in the neurosciences are revealing that the animal brain is more flexible or plastic than previously thought. Psychologists are demonstrating that the contexts of experience can be manipulated to reveal that human infants are capable of learning a broader range of mental skills than was previously assumed possible. These emerging scientific fronts in machine and human learning create the possibility to scale up our understanding of mental development in complementary ways.

On the one hand, psychologists historically concerned with the development of the mind, have produced important findings, but computational studies of mental development, which are demanded by better understanding of both human and machine intelligences, have been lacking. On the other hand, traditional research paradigms in machine learning have fruitfully been informed by models of human learning. However, existing online learning techniques typically applied to robot learning (for example, Hexmoor, Meeden, and Murphy [1997]) differ fundamentally from human learning. Online root learning using robot sensors is not equivalent to autonomous mental development in robots, nor should mental development be viewed simply as an incremental learning system that grows from small to big in terms of its memory size. Such systems already exist (for example, systems that use neural network techniques).

There is a need, therefore, for increased studies in computational autonomous mental development (CAMD) that are of interest to both machine and human intelligence researchers. The workshop was motivated generally by the need to form a new alliance among the disciplines concerned with mental development in machines or humans. The workshop provided a forum for articulating a new research direction and basis for scholarly collaboration across related disciplines. This was the first workshop dedicated solely to this new direction, CAMD. The workshop was intended to be a catalyst for creating a new research community interested in advancing the science within this new paradigm.

What are the central issues of CAMD by robots and animals? What does neuroscience tell us about mental development? What computational studies for mental development are needed in neuroscience and psychology? How does a robot develop its cognitive and behavioral skills autonomously? Answers to these and other related questions have the potential to advance AI. Further, they might also affect human life in ways that have practical consequences. For example, developmental machines can be created and used to interact with humans and help them to learn. Intelligent developmental robots can also be used to test models of human mental development and learning. The power to computationally model human mental development and learning well enough to predict particular behavioral outcomes would reflect better understanding of the processes that result in different human behavioral outcomes. As a result, it might be possible to manipulate the processes well enough to create more sophisticated learning tools or procedures that enhance normal and abnormal human learning.

Planning and Organization

A grant jointly awarded by the National Science Foundation and the Defense Advanced Research Projects Agency enabled us (Co-PIs) to plan and conduct such a workshop. Participants attended by invitation only. They represented multiple disciplines that included developmental and cognitive psychology, neuroscience, robotics, and AI. The accepted invited submissions appeared in the WDL proceedings as related working papers at the conference and can be obtained on request. In the month preceding the conference (March 2000), a preworkshop e-mail discussion among potential participants was initiated to stimulate discussion on the workshop themes.[1]

Workshop Schedule

The workshop was held on 5 to 7 April 2000 at Michigan State University's Kellogg Conference and Hotel Center. The participants included more than 30 invited scholars, most of whom resided in the United States. As planned, the first two days of the workshop were devoted to focal scholarly exchanges in four designated research areas: (1) developmental psychology, (2) computational modeling, (3) neuroscience, and (4) robotics. The third and final day focused on future directions. Both morning and afternoon sessions during the first two days were structured to include (1) an overview talk that identified relevant major issues in a field, (2) short respondent papers that addressed related themes, and (3) an hour-long spontaneous discussion period. On the third day, an hour-long discussion followed several short overview talks. The last session was devoted to open discussion from the floor.

Presentations and Discussions

The workshop discussion included neuroscientists whose studies have provided a window for better understanding brain development. Neuroscientist Mriganka Sur presented research in which the optic nerves originating from the eyes were rewired into the auditory cortex of an animal (a ferret) early in life (von Melchner, Palls, and Sur 2000). They found that the rewired auditory cortex gradually takes on a representation that is normally observed in the visual cortex. Further, the animals successfully learned to perform visual tasks using the auditory cortex. In other words, the rewired ferrets can see using the brain zone that is normally assigned to sound. This discovery suggests that the cortex is governed by self-organizing mechanisms that can derive their representation and even partial architecture from either visual or auditory input signals. An interesting question to ask is, Is there a set of general developmental mechanisms that is applicable to vision, audition, and touch?

If the genetically assigned hearing zone of a newborn can be developed into a seeing zone, how rigid is the brain of an adult? Neuroscientist Michael M. Merzenich presented evidence for great plasticity in the adult brain of nonhuman primates (Wang et al. 1995). He and his colleagues have shown that finger skin areas from which a neuron in the somatic cortex receives sensory signals (called the *receptive field* of the neuron) can change according to sensory experience. If multiple fingers of the adult monkey receive consistent synchronized pulse stimuli from a cross-finger bar for several days, the receptive field changes drastically, from covering only a single finger in typical cases to covering multiple fingers. This study seems to demonstrate that the brain's self-organizing program still automatically selects the source of sensory input within a candidate area according to the statistical properties of the actual sensory signal that is received. How plastic should a developmental robot be?

Developmental psychologists, Esther Thelen (Thelen et al. 2002), Rachel Clifton, Neil Berthier, and Kurt Fischer discussed studies of early human cognition from a dynamic systems theory perspective. This view grounds mental learning and developing in the physical contexts of sensorimotor body interactions with the environment. What is known is that normal mental development is context dependent and inseparable from the biological, social, and psychological contexts of experience. Psychologists Kim Plunkett, James McClelland, and Nestor Schmajuk presented their studies and theoretical perspectives on mental development from a neural network point of view. Plunkett, a coauthor of a recent book entitled *Rethinking Innateness* (Elman et al. 1997), summarized studies about the plasticity of the cortex for generating representations according to the signals it receives. McClelland suggested that Hebbian learning mechanisms might play an important role in cortical development. Schmajuk presented computational models for autonomous animal learning based on classical and instrumental conditioning, attention, and model generation about the environment. He proposed that the ideas of these learning models could also be used for robots.

Computer scientists and robot researchers Tomaso Poggio, Roderic Grupen, Maja Mataric, and Brian Scassellati presented their work on computational learning based on network, Markov decision process, and behavior-based approaches. Although they have not yet reached a fully autonomous developmental system with incrementally self-generated representation, these studies are motivated by human animal development or learning. For example, the work of the COG Project by Rodney Brooks and other coworkers (Brooks et al. 1999), presented by Brian Scassellati at the workshop, is motivated by child cognitive development in infancy. Alex Pentland discussed work that has associated video images of an object with a synchronized voice (that is, pronounced the verbal name of an object) (Roy, Schiele, and Pentland 1999). Computer scientist Stephen Levinson's past work focused on machine recognition of speech and language modeling. He suggested that a breakthrough in speech recognition is unlikely without a fundamental change of methodology. He has started working on a project that enables a mobile robot to develop its capability through interaction with the environment.

Early examples of fully autonomous developmental robots include the DARWIN V robot at

The Neurosciences Institute in San Diego, California, and the SAIL robot at Michigan State University, developed independently about the same time but with very different goals. Gerald Edelman and Olaf Sporns presented their work on DARWIN V (Almassy, Edelman, and Sporns 1998), whose goal was to provide a concrete example of how the properties of more complex and realistic neural circuits are determined by the behavioral and environmental interactions of an autonomous device. DARWIN V has been tested for the development of generalization behaviors in response to visual stimuli at different positions and orientations (visual invariance learning). It has been shown that through real-time interaction with this environment, the robot can develop more complex behaviors from simpler ones.

SAIL was designed as an engineering test bed for developmental programs that are meant for scaling up to complex cognitive and behavioral capabilities in uncontrolled environments (Weng 1998; Weng et al. 1999). Juyang Weng presented the design principles of the SAIL-3 developmental program with a goal to autonomously develop some cognitive skills for vision, speech, and touch through real-time interaction with the environment. The SAIL developmental program does not require two separate modes: (1) a training mode and (2) a performance mode. The SAIL robot learns while it performs under the same operational mode so that it can truly scale up. Weng called it *autonomous animallike learning*.

Technically, Weng explained how a real-time speed with a large memory, as well as autonomous derivation of high-dimensional discriminating features, is achieved by an incremental high-dimensional mapping engine called *hierarchical discriminate regression* (Hwang and Weng 2000). He suggested that a true developmental program must autonomously and incrementally generate representations online from sensory and effector signals. Such representations include filters, feature spaces, clusters in feature spaces, feature hierarchies (architecture), internal effectors (for example, for internal attention), and the association of response in self-generated context representation with the corresponding actions and their expected values.

Future Research Issues

The workshop dedicated a half day to identifying future research issues. Christopher Brown, Sporns, Weng, and Stan Franklin presented their related work, research challenges, and possible future research topics.

Some researchers pointed out that to truly scale up cognitive capabilities, fully autonomous online, real-time mental developmental robots are necessary, which, in turn, requires a fundamental paradigm change (Weng et al. 2001). Such robots are also necessary for integration of a wide variety of perceptual and behavioral capabilities. A fundamental characteristic that distinguishes a developmental program from all other programs is that it is not task specific—the tasks that the robot will learn are not known to the programmer at programming time.

Some researchers hold the view that the task-nonspecific nature of mental development should make the studies of intelligence and realizations of intelligence easier than the traditional task-specific approaches, which is true both for human subjects (neuroscience and psychology) and machine subjects (AI and robotics). From the computational view of mental development, however, the research issues center on how to autonomously generate internal representations from sensory and effector signals.

This new research area opens up a series of very interesting and yet manageable new research topics for fields that study either human or machine subjects. Some of the tractable research topics suggested at the workshop include:

Representation from sensory signals: This topic includes schemes for automatic derivation of mental representations from sensory signals that are sensed from the environment and the body.

Representation from effector signals: This topic includes schemes for automatic derivation of representation from effector signals, available from practice experience.

Automatic derivation of receptive fields in both the classic and nonclassic sense: This topic explores how later processing elements by the brain can group outputs from earlier processing elements or sensory elements.

Long-term memory growth, self-organization, and retrieval for high-dimensional neural signal vectors: This type of real-time memory engine serves the role of a cortex. It must be scalable (for example, logarithmic time complexity in memory size) so that it is always real time when the memory grows.

Working memory formation and self-organization for high-dimensional neural signal vectors: The working memory can include short-term sensory memory and system states as currently attended context.

Developmental mechanisms for mediation of conscious and unconscious behaviors: This topic covers mechanisms for mediation among higher- and lower-level behaviors, such as learned behaviors, learned and innate emotional behaviors, and reflexes.

Mechanisms for developing internal behaviors: This topic includes those mechanisms that operate on internal nervous components, including selective attention. This subject includes both developmental mechanisms and training strategies for humans and robots.

Attention-directed time warping from continuous states: The time warping issue concerns the time inconsistency between different instances of experience, with the goal of both discrimination and generalization.

Autonomous action imitation and self-improvement: This topic includes developmental mechanisms underlying an improved behavioral pattern that results from individual online instances of related experience.

Mechanisms for communicative learning and autonomous thinking: Communicative learning refers to learning directly from interpreting or associating previously internalized meanings or representations with linguistic forms delivered in one or more modality codes (that is, auditory speech, visual print, or tactile braille), similar to how children learn when they attend classes. This type of learning, in turn, requires language acquisition in a physically grounded way. These mechanisms interact with the processes underlying the development of thinking, which is responsible for categorizing, planning, decision making, and problem solving.

A developmental program for machines creates many interesting technical problems that can be addressed by researchers who work with computerized algorithms. Indeed, it is a great challenge to design a scalable, real-time developmental program, namely, one that is open ended; non–task specific; and, therefore, capable of learning from its own history of experiences in the environment, including interactions with humans.

Final Remarks

This article reports on the WDL workshop and the new cross-disciplinary initiative it represents. The workshop raised many research issues that are of fundamental importance to AI, neuroscience, psychology, and robotics. Mental development seems to be a common research subject where these fields can find common ground. The area of CAMD raises many questions that await investigation by researchers in related fields.

In the closing workshop discussion, participants recommended creating a web site repository on development and learning,[2] preparing a white paper (now published in *Science* [Weng et al. 2001]), and establishing a regular conference series. The Second International Conference on Development and Learning (ICDL'02) will be held at MIT on 12 to 15 June 2002.[3]

Notes

1. These discussion items are still available under the E-mail Forum at the WDL web site: www.cse.msu.edu/dl/.

2. It is now available online at www.mentaldev.org/.

3. See the conference web site at www.egr.msu.edu/icdl02/.

References

Almassy, N.; Edelman, G. M.; and Sporns, O. 1998. Behavioral Constraints in the Development of Neural Properties: A Cortical Model Embedded in a Real-World Device. *Cerebral Cortex* 8(4): 346–361.

Brooks, R. A.; Breazeal, C.; Marjanovic, M.; Scassellati, B.; and Williamson, M. M. 1999. The COG Project: Building a Humanoid Robot. In Computation for Metaphors, Analogy, and Agents, ed C. L. Nehaniv. Lecture Notes in Artificial Intelligence 1562. New York: Springer Verlag.

Elman, J. L.; Bates, E. A.; Johnson, M. H.; Karmiloff-Smith, A.; Parisi, D.; and Plunkett, K. 1997. *Rethinking Innateness: A Connectionist Perspective on Development.* Cambridge, Mass.: MIT Press.

Hexmoor, H.; Meeden, L.; and Murphy, R. R. 1997. Is Robot Learning a New Subfield? The Robolearn-96 Workshop. *AI Magazine* 18(4): 149–152.

Hwang, W. S., and Weng, J. 2000. Hierarchical Discriminate Regression. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 22(11): 1277–1293.

Roy, D.; Schiele, B.; and Pentland, A. 1999. Learning Audio-Visual Associations Using Mutual Information. Paper presented at the International Conference on Computer Vision, Workshop on Integrating Speech and Image Understanding, 21 September, Corfu, Greece.

Thelen, E.; Schoner, G.; Scheier, C.; and Smith, L. B. 2001. The Dynamics of Embodiment: A Field Theory of Infant Perseverative Reaching. *Behavioral and Brain Sciences* 24:1–34.

Von Melchner, L.; Pallas, S. L.; and Sur, M. 2000. Visual Behavior Mediated by Retinal Projections Directed to the Auditory Pathway. *Nature* 404(6750): 871–876.

Wang, X.; Merzenich, M. M.; Sameshima, K.; and Jenkins, W. M. 1995. Remodeling of Hand Representation in Adult Cortex Determined by Timing of Tactile Stimulation. *Nature* 378:71–75.

Weng, J. 1998. Learning in Computer Vision and Beyond: Development. In *Visual Communication and Image Processing,* eds. C. W. Chen and Y. Q. Zhang, 431–487. New York: Marcel Dekker.

Weng, J.; Hwang, W. S.; Zhang, Y.; and Evans, C. 1999. Developmental Robots: Theory, Method, and Experimental Results. Paper presented at the Second International Symposium on Humanoid Robots, 8–9 October, Tokyo, Japan.

Weng, J.; McClelland, J.; Pentland, A.; Sporns, O.; Stockman, I.; Sur, M.; and Thelan, E. 2001. Autonomous Mental Development by Robots and Animals. *Science* 291(5504): 599–600.

Juyang Weng (B.Sc., computer science, Fudan, Shanghai, 1982, M.Sc., and Ph.D., computer science, University of Illinois at Urbana-Champaign, IL, 1985 and 1989, respectively) is an associate professor in the Department of Computer Science and Engineering, Michigan State University. His research interests include mental development; computer vision; autonomous mobile robots; humanoid robots; autonomous navigation; and human-machine multimodal interactions using vision, audition, touch, speech, gesture, and actions. His e-mail address is weng@cse.msu.edu.

Ida Stockman (M.S., speech-language pathology, University of Iowa, 1965, and Ph.D., speech-language pathology and child development, Pennsylvania State University, State College, PA, 1971) is a professor in the Departments of Audiology and Speech Sciences and Counseling, Educational Psychology, and Special Education, Michigan State University. Her research interests include human development, particularly the role that culture and different sensory modalities play in the perception and acquisition of spoken language and nonverbal cognition in children with typical and atypical developmental histories. Her e-mail address is stockma1@msu.edu.