IEEE CIS

AMD Newsletter

Robotics
Psychology
Neuroscience
Machine Intelligence

The Newsletter of the Autonomous Mental Development Technical Committee

VOLUME 2, NUMBER 1 | ISSN 1550-1914 | APRIL 2005

Message From The Editor

Welcome to the first issue of the AMD TC Newsletter in 2005. The highlight of this issue is the dialog column guest-moderated by Olaf Sporns: “What are Potential Principles for the Autonomous Development of Value and Motivation?” I would like to thank Olaf and all contributing authors of this column for their insights. The AMD TC Newsletter welcomes more contributions from the AMD TC on research topics interesting to the community.

- Yilu Zhang, Editor of the AMD Newsletter

Committee News

- October, 2004: The AMDTC and ICDL business meeting was held at the 3rd ICDL (2004) at La Jolla, CA. A resolution about changing the name of the ICDL Advisory Board to the ICDL Governing Board (GB) and a procedure to elect the ICDL GB were passed. Each AMD task force presented its past work and future plan.

- February, 2005: With Jay McClelland as the ICDL GB Interim Chair and Jochen Triesch as the communication co-Chair, the election of two new members of ICDL GB and the Chair of ICDL GB has taken place. Hideki Kozima and Linda Smith were elected to the board by the membership. Jay McClelland, Sandy Pentland, Terry Sejnowski, and John Weng are continuing from the ICDL Interim GB for continuity and will rotate off in 1-2 years. John Weng was elected to serve as the first ICDL GB Chair, replacing Jay McClelland, who served as Interim Chair while the election of new members was conducted. Jochen Triesch will continue to serve as communications co-Chair of the GB, serving a three-year term which began in the fall of 2004.

Dialog Column

Dialog: What are Potential Principles for the Autonomous Development of Value and Motivation?

Olaf Sporns, Department of Psychology, Cognitive Science and Neuroscience, Indiana University, Bloomington, IN 47405, USA

One of the goals of research in autonomous mental development is to formulate principles that enable adaptive and generalizable behavior in humans and robots. Clearly, internal evaluations and motivations, in the context of environmental stimuli, action, goals, and plans, play important roles in the development of behavior. As we attempt to build systems that exhibit autonomous development, what have we learned about the role of value and motivation in behavioral development?

Value can be seen as imposing biases on the outcome of interactions between a behaving system and its environment, thus providing a way to evaluate recent behavior and perception. Value systems, as part of a robot's control architecture or an organism's brain, then exert effects that change the propensity or likelihood of such behaviors to occur in the future, in a manner similar to reinforcement learning. Signals produced by value systems represent global evaluations that are internally derived (not externally supplied), and are communicated to widespread regions of the brain or neural network.

Dialog Column

Stimuli that trigger value systems (i.e. that are rewarding, noxious, novel, or violate expectations) have a greater impact on learning and development. Value may be triggered by touching or manipulating an object, the appearance of high-contrast visual features such as blobs or edges, or by stimulus motion [3]. Evidence from developmental psychology suggests that in human infants, sensitivity to highly salient sensory inputs emerges very early in postnatal development and can lead to behavioral changes towards behavioral patterns that are more strongly associated with the recurrence of value. A very important issue concerns the necessity for value systems to be themselves capable of adaptation and plasticity. Clearly, only very few stimuli are intrinsically (innately) coupled to value - most relationships between external stimuli and internal evaluations must be the result of actual experience, as they cannot be encoded by evolution.

In the brain, modulatory neurotransmitter systems, including dopamine, serotonin, and acetylcholine, possess many of the structural and functional properties of value systems, including their capacity to acquire new response characteristics in the course of experience. Such systems may play important roles in neural and cognitive development by influencing the magnitude and direction of synaptic plasticity throughout the brain, and thus generating lasting behavioral and representational change [1]. A developing robot, built on biological principles, should therefore emulate the pivotal role of neuromodulation in real brains [2].

I pose the following questions for discussion: Why is a value or motivational system necessary or important for a developing human/machine? What specific mechanisms exist to enable value systems to become and remain plastic, adaptive, and self-organizing throughout development? Can we classify different kinds of value systems by virtue of their architecture or mechanisms? How can such systems be implemented in autonomous robots and agents?

References:
[1] K. Doya, "Metalearning and neuromodulation," *Neural Networks* vol.15, pp. 495-506, 2002.
[2] O. Sporns and W.H. Alexander, "Neuromodulation and plasticity in an autonomous robot." *Neural Networks* vol.15, pp. 761-774, 2002.
[3] E. Thelen and L.B. Smith, *A Dynamic Systems Approach to the Development of Cognition and Action*. Cambridge, MA: MIT Press, 1994.

Implementing Autonomous Development of Value and Motivation via Feelings and Emotions

Stan Franklin, Computer Science Division and the Institute for Intelligent Systems, University of Memphis, Memphis, TN 38152, USA

Olaf Sporns asks, "Why is a value or motivational system necessary or important for a developing human/machine?" Every autonomous agent, be it human, animal, robot, or software agent, must iteratively and continually sense its environment and act on it in pursuit of its own agenda [3]. Whether that agenda is derived from drives built in by the designer of an artificial agent or evolved into a biological agent, there must be motivation (a motive generator) for the agent's choosing its next action [6]. By estimating the value of its outcome in relation to current goals, an agent can bias motivation regarding a particular action. In biological agents such motive generators are typically implemented via feelings and emotions. (By an emotion, I mean a feeling with cognitive content.) Hunger often leads to feeding activity, fear frequently to escape activity. Recently, artificial emotions have been used to motivate artificial agents [1] [5].

Sporns also asks, "What specific mechanisms exist to enable value systems to become and remain plastic, adaptive, and self-organizing throughout development?" and "How can such systems be implemented in autonomous robots and agents?" In my own IDA technology, a conceptual and computational model of cognition, we are implementing

Dialog Column

motivation and biasing several forms of learning via feelings and emotions [4]. These forms of learning, particularly perceptual learning [2] and procedural learning, enable an artificial autonomous agent, during its developmental periods and thereafter, to associate various feelings with situations it perceives in its environment, and with the consequences of its actions. Elements of a situation, possibly of the result of an action that results in pain, may be associated with fear during perceptual learning.

Perceptual and procedural memory are both implemented in the IDA technology by means of a semantic net with complex activation passing. The value system is implemented implicitly in the consciousness mechanism and the action selection mechanism (behavior net).

References:

[1] L.D. Canamero, "Designing Emotions for Activity Selection in Autonomous Agents," in *Emotions in Humans and Artifacts*, ed. R. Trappl, P. Petta, and S. Payr. Cambridge, MA: MIT Press, 2003.

[2] S. Franklin, "Perceptual Memory and Learning: Recognizing, Categorizing, and Relating," presented at American Association for Artifical Intelligence (AAAI) Symposium on Developmental Robotics, Stanford University, Palo Alto CA, USA. March 21-23, 2005. (http://www.cs.brynmawr.edu/DevRob05/schedule/)

[3] S. Franklin and A. C. Graesser, "Is it an Agent, or just a Program?: A Taxonomy for Autonomous Agents," in *Intelligent Agents III*. Berlin: Springer Verlag, 1997.

[4] S. Franklin, and L. McCauley, "Feelings and Emotions as Motivators and Learning Facilitators." *Architectures for Modeling Emotion: Cross-Disciplinary Foundations*, AAAI 2004 Spring Symposium Series; American Association for Artificial Intelligence; Stanford University, Palo Alto, California, USA; March 22-24, 2004. (http://csrg.cs.memphis.edu/csrg/html/papers.html)

[5] R. Picard, *Affective Computing*. Cambridge, MA: MIT Press, 1997.

[6] A. Sloman, "Motives, mechanisms, emotions." *Cognition* vol.1 pp. 217-234, 1987.

Reply to "What are Potential Principles for the Autonomous Development of Value and Motivation?"

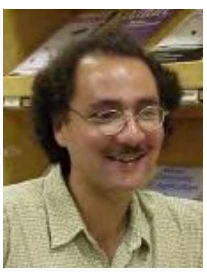

Jean-Marc Fellous, Department of Biomedical Engineering and Center for Cognitive Neuroscience, Duke University, Durham, NC 27708, USA

To be successful, an autonomous agent (software, robot or animal) needs to be adaptive and motivated. Olaf Sporns raises a number of interesting points that are echoed by a careful study of the nervous system. Unfortunately, little has been done at this point in using what is known about the evaluative system of the mammalian brain when building autonomous motivated devices. Two main aspects of this system (there are certainly many others) seem particularly important to keep in mind when thinking about its possible implementation.

First and foremost the evaluative system of the brain is decentralized. There is no single 'module' in charge of attributing a value to a percept or an action. While it is possible to affect the overall evaluation capabilities of an organism by selectively inactivating certain brain areas, not all evaluative capabilities are affected, and deficits other than evaluative ones are also observed. This fundamental aspect of evaluative function is rarely observed in artificial autonomous agents. I would argue that because of this, their evaluative abilities are limited and cannot develop efficiently. The limitation in capacity comes from the unavoidably limited amount of processing power available to a single module, compared to a distributed set. The limitation in development comes from the large amount of interdependencies inherent to compact systems whereby the modification of one aspect of the evaluation affects many other aspects of other evaluations, and is therefore unstable unless it is tightly regulated, or even discouraged.

Dialog Column

A second aspect is that the brain's evaluative system works at different levels, from stereotypical reflex-like levels such as playing piano, to higher 'cognitive' levels such as planning a set of behaviors to achieve a goal. Each level requires a different amount of expertise and sophistication in its associations between particular values and particular percepts or actions. The mechanisms for evaluations might be fundamentally different at each of these levels, allowing for more efficiency in stimulus representations and processing. This organizational principle suggests therefore that the implementations of evaluative functions in autonomous agents would perhaps gain from multi-level design with values being attributed at the hardware, operating system level or user-software levels simultaneously.

A multi-level, distributed approach to constructing evaluative and motivated autonomous agents seems to be essential to achieve animal-like motivated behaviors that are based on an evaluative system [1]. The major challenge in developing such a system is to find ways of making the different components of the evaluative system communicate efficiently, at different levels of abstraction, a task that is still as of today central to Artificial Intelligence research.

References:

[1] M.A. Arbib and J-M. Fellous, "Emotions: from brain to robot." *Trends in Cognitive Sciences* vol. 8, pp. 554-561, 2004.

Adaptive internal value systems for developmental robots

Pierre-Yves Oudeyer and Frédéric Kaplan, Sony Computer Science Laboratory Paris, 6, Rue Amyot F-75005, Paris, France

Autonomous mental development requires a system which attributes internal values, or rewards, to sensorimotor experiences. Searching to maximize these internal rewards is a way to drive the development of a robot without the need for externally provided rewards. These internal rewards can possibly be generated when particular stimuli are encountered, such as the vision of a face, the touching of an object, or the sensation of movement. This kind of internal reward can be useful for a robot, but because they are fixed in advance, they limit severely its development. In particular, once the robot has learnt robust strategies for finding the corresponding rewarding situations (e.g. finding a face), nothing will attract the robot towards new activities which might push further its development. Its development stops there.

Therefore, there is a need for adaptive internal value systems for which sensorimotor experiences that produce positive rewards evolve with time. Furthermore, this evolution should be such that it continually and efficiently drives the robot towards the acquisition of new skills, with increasing structure and complexity. Indeed, not all adaptive value systems lead to a structuring of the behavior of the robot (e.g. random attribution of values to each situation).

One possibility for designing adaptive internal value systems is to rely on abstract and intrinsic properties of the sensorimotor flow. This boils down to measures that characterize the evolution of the states of sensorimotor channels independently of the "meaning" of those channels. This means that these measures should be useable without knowing if they are applied for example to the flow of auditory sensors states, or of visual sensors states, or of joint speed states, or of the coupling of all these states.

There are already in the literature several kinds of measures which fulfill these requirements, and which can be used to generate internal rewards. Moreover, some researchers have developed examples of implementation for each of them:

1) synchronicity: this characterizes the high temporal correlational information between the evolution of various sensorimotor channels (e.g. [4]).
2) novelty: this characterizes the low probability of the occurrence of certain sensorimotor situations, or of transitions between sensorimotor situations according to an internal model of their distribution (e.g. [2]).

Dialog Column

3) prediction error: this characterizes the high errors made by an internal machine predictor trying to anticipate the next sensory state (e.g. [5]).
4) information gain: this characterizes the decrease of the entropy of an internal model of the distributions of sensorimotor situations (or of their transitions). Theoretical properties were developed by [1] in the context of optimal experiment design. This measure is computationally expensive and has actually never been implemented on an autonomous agent.
5) learning progress: this characterizes improvement in the performances of an internal machine predictor trying to anticipate the next sensory state (e.g. [3]).

The challenge for the next years will to be show experimentally how such forms of intrinsic value systems lead to complex organized developmental sequences.

References:
[1] V. Fedorov, *Theory of Optimal Experiment.* Academic Press, 1972.
[2] X. Huang and J. Weng, "Novelty and Reinforcement Learning in the Value System of Developmental Robots." In Prince, Christopher G. and Demiris, Yiannis and Marom, Yuval and Kozima, Hideki and Balkenius, Christian, Eds. *Proceedings Second International Workshop on Epigenetic Robotics: Modeling Cognitive Development in Robotic Systems* 94, Edinburgh, Scotland, 2002, pp. 47-55.
[3] P-Y Oudeyer, F. Kaplan V. Hafner, and A. Whyte, "The Playground Experiment: Task-Independent Development of a Curious Robot," in *Proceedings of the AAAI Spring Symposium Workshop on Developmental Robotics*, 2005. http://playground.csl.sony.fr
[4] C. Prince, G. Hollich N. Helder, E. Mislivec, A. Reddy, S. Salunke, and N. Memon, "Taking Synchrony Seriously: A Perceptual-Level Model of Infant Synchrony Detection," in *Proceedings of the Fourth Workshop on Epigenetic Robotics*, Lund University Cognitive Studies, 2004, pp. 89-96.
[5] S. Thrun, "Exploration in active learning." In *Handbook of Brain and Cognitive Science*. M. Arbib, ed. Cambridge, MA: MIT Press, 1995.

Motivation-free and Motivational Systems

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

We have seen a series of approaches: direct programming (e.g., inborn reflexes and many knowledge-based systems), supervised learning (e.g., holding hands to teach a child to draw and most pattern recognition systems), reinforcement learning (e.g., instrumental conditioning and Q-learning capable robots) and autonomous mental development (e.g., human children and developmental robots). From earlier to later approaches, this series shows an increase in the capability of the system to develop complicated cognitive and behavioral capabilities in unknown, unpredictable, and complex environments and in being unbounded by the limit of human task-specific programming. Direct programming and supervised learning are free of a motivational system. Reinforcement learning [2] has a motivational system that has a fixed desired value (e.g., physical pleasure or Q-value). Among other things, the motivational system enables the machine to acquire complex capabilities (e.g., avoiding situations that cause pain) *more autonomously* (e.g., without explicitly coding which actions are painful).

Autonomous mental development goes beyond reinforcement learning. Reinforcement can create a pain evader. However, a machine or a human with a fully developed motivational system is able to generate goals (e.g., helping humans) and sub-goals (e.g., daily duties) autonomously and take much pain to achieve the goals (i.e., resolving conflicting short and long term goals). Further, it is able to be taught about what is right and what is wrong from verbal communications while no pain or pleasure takes place. Thus, a developed motivational system has *adaptive desired values*.

Dialog Column

Therefore, all systems fall into four categories: (1) learning-free systems, (2) learning systems without a motivational system (e.g., pattern recognition), (3) learning systems with fixed desired values (e.g., Q-learning), and (4) learning systems with adaptive desired values (changed by experience). The first two are motivation-free systems and the latter two are motivational systems.

In the past, physical rewards and novelty have been largely studied separately. Xiao & Weng [1] have proposed a unified developmental motivational system for the SAIL robot in the context of autonomous development using high-dimensional stimuli. It has three components: aversive stimuli (negative reward), appetitive stimuli (positive reward) and novelty (deviation from experience-based expectation), with earlier ones having a higher priority (e.g., a strong pain is more urgent than looking for novelty). It realizes a variety of animal learning, including non-associative learning (e.g., sensitization and habituation) and associative learning (e.g., classical conditioning and instrumental conditioning). For example, the SAIL robot turned its head away (changed value) after looking at a scene for a while (no aversive or appetitive stimuli took place). This motivational system goes beyond reinforcement learning since experience without rewards also changed the value. Further work is needed to demonstrate whether this developmental model of motivational system is able to develop a type (4) system for complex concepts of adaptive values, such as assisting humans.

References:
[1] X. Huang and J. Weng, "Novelty and Reinforcement Learning in the Value System of Developmental Robots," in *Proc. Second Int'l Workshop on Epigenetic Robotics*, Edinburgh, Scotland, August 10-11, 2002.
[2] R. S. Sutton and A. Barto, *Reinforcement Learning*, Cambridge, MA: MIT Press, 1998.

Reply and Summary:
"What are Potential Principles for the Autonomous Development of Value and Motivation?"

Olaf Sporns
Department of Psychology, Cognitive Science and Neuroscience, Indiana University, Bloomington, IN 47405.

Adaptive motivational systems are central components of cognitive and neural architectures capable of autonomous mental development (AMD). There seems to be unanimous support for this notion among the contributors to this dialog column. All authors note that effective and powerful value systems must be continually modifiable by experience. Value systems that encode a single, fixed stimulus dimension can produce adaptive behavior, but do not enable open-ended continuous development, pushing the system towards higher levels of behavioral complexity (Oudeyer and Kaplan). Fellous raises the extremely important point that evaluative systems need to be "multi-level," ranging from reflex-like processes to higher cognitive function. A single, centralized value system is unlikely to be able to carry out operations covering this entire functional range. Communication and integration across the multiple components of such a multi-level evaluative system is a major issue for future AMD research. Weng provides a very useful taxonomy of adaptive motivational systems and notes that AMD goes beyond reinforcement learning, in that the definition of the "desired state" must be a function of the actual experience of the autonomous system. Weng stresses the importance of novelty, while Oudeyer and Kaplan cite increasing structure and complexity as key driving forces for autonomous development. Franklin, Kaplan, and Weng have designed and studied architectures that implement several of these principles.

The degree to which these different viewpoints are in basic agreement with one another is encouraging. However, a number of open questions remain. While there is some agreement on general principles, or at least "desirable design features" of adaptive motivational systems, specific implementations and mechanisms are still in short supply. How can principles or mechanisms of autonomous motivational systems be validated, tested, and consolidated into a single theoretical framework? Here, I think the empirical study of neural and cognitive processes in developing brains and organisms (especially humans) plays an important role. For example, several basic design principles of mammalian

Dialog Column

neural circuits underlying value, motivation, and emotion have been discovered, but relatively little has been accomplished in implementing these principles in autonomous systems, agents, or robots (a point also argued by Fellous). Implementing specific architectures in artificial behaving systems may have a dual benefit: it could help in understanding their biological function, while also providing us with devices that have a broad range of autonomous behaviors and capabilities.

I want to raise one more issue, which relates to the connection between the agent and the environment. In my view, one of the engines of continuous development in an embodied agent is its dynamic coupling to an environment that is partly shaped and structured by the agent's own behavior. In other words, we need to consider the agent and its physical/social environment as a single coupled system, rather than as separate domains, one of which develops (the agent) while the other remains unchanged (the environment). Evaluations and motivations give direction to behavior and thus impose structure (in a statistical sense) on sensory inputs that are received and processed by the agent's brain [1]. This, in turn, induces changes in the agent's control architecture, including its motivational system. A possible hypothesis is that the rate of change of environmental statistics due to the actions of the agent may serve as a "cost function" for driving open-ended developmental change.

I thank all contributors for their thoughtful comments.

References:

[1] M. Lungarella, T. Pegors, D. Bulwinkle, and O. Sporns, "Methods for quantifying the informational structure of sensory and motor data." *Neuroinformatics* (forthcoming).

Conference Reports


ICDL'04 - Here Comes the Rain Again

Jochen Triesch, Department of Cognitive Science, University of California, San Diego, CA 92093-0515, USA

Tony Jebara, Department of Computer Science, Columbia University, New York, NY 10027, USA

Who would associate San Diego with pouring rain? Not many of you, I assume. But then again, some of you – those who attended ICDL'04 – may have had more than enough evidence that it does rain in southern California. In fact, sometimes it even pours!

ICDL'04, the third International Conference on Development and Learning, was held in sunny San Diego on the rainy days of October 20-22, 2004. It was hosted by the Salk Institute for Biological Studies, who let us use their premises free of charge (thank you!). In line with the earlier ICDL conferences, the San Diego meeting had a single track format. Over 30 regular talks and over 40 posters were presented during the three days. In addition, we were fortunate to have a truly remarkable set of invited talks by distinguished colleagues such as (in order of appearance) Terry Sejnowski, Dana Ballard, Jay McCLelland, Eric Courchesne, Bill Greenough, and Karen Dobkins. For the first time the conference also had a special theme: "Developing Social Brains." All the papers and abstracts have been published online and can be downloaded from the conference web site (http://www.icdl.cc).

The meeting was kicked off on Wednesday morning by some rain and a set of four tutorials given by leading experts in their respective fields. Leslie Carver and Gedeon Deak represented Developmental Psychology; from Sanjoy Dasgupta we learned about Machine Learning; Irene Merzlyak instilled Neuroscience knowledge into our synapses; and Juyang Weng developed our understanding of Autonomous Mental Development. That afternoon, accompanied by some rain, the main program started with two oral sessions on "Attention and Learning in Social Systems."

Wednesday night's poster session sported presentations in the areas of "Social Robots" and "Social Systems." Thursday was devoted to some rain and oral sessions on "Reinforcement and Neuromodulation" and "Development and Language." The evening poster session featured research on "Perceptual Learning and Development," and "Developmental Disorders" among many other things. Friday saw (you must have guessed it) more rain and sessions on "Brain, Emotion, and Social Dynamics," "Perceptual Learning and Development," and the conference concluded

Conference Reports

with a session on "Binding and Modularity". Other highlights of the conference included the best paper award, which was given to Andrew Lovett and Brian Scassellati for their contribution "Using a Robot to Reexamine Looking Time Experiments" (congratulations!), and Minoru Asada's bid to organize ICDL'05 (http://www.icdl05.org/) in Osaka, for which we are very grateful!

By the third day of the conference, the "lucky" ICDL'04 attendants had witnessed a fair share of something quite remarkable: in October 2004, San Diego saw 4.98 inches of rainfall - the highest mark in over 150 years. But despite the terrible weather throughout the conference, we all had a wonderful time. And on Saturday, finally, when a few of us went on a post-conference hike in the beautiful Torrey Pines State Park, believe it or not, the San Diego sun was shining again - as it (almost) always does in sunny San Diego.

Call For Participation

International Conference
on Fuzzy Systems (FUZZ -IEEE 2005)
May 22-25, 2005, Reno, Nevada, USA
http://www.fuzzieee2005.org/

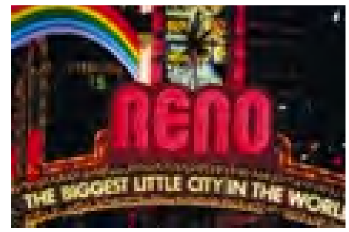

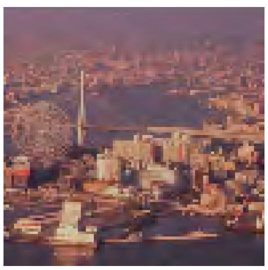

The Fourth International Conference
on Development and Learning (ICDL '05)
"From Interaction To Cognition"
July 19-21, 2005, Osaka, Japan
http://www.icdl05.org/

International Joint Conference
on Neural Networks (IJCNN 2005)
July 31-August 4, 2005, Montreal, Canada
http://faculty.uwb.edu/ijcnn05/

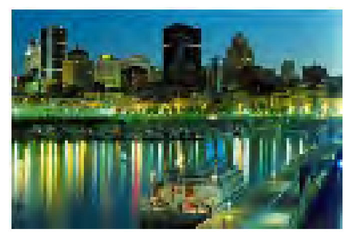

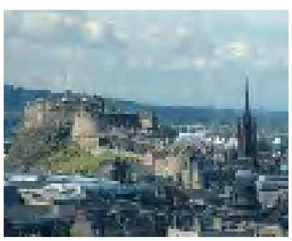

IEEE Congress on Evolutionary Computation
September 2-5, 2005, Edinburgh, UK
http://www.dcs.ex.ac.uk/~dwcorne/cec2005/

Glossary

Motivational system: Motivational system is a term that has emerged from the disciplines of neural science and psychology. In psychology, many researchers also use the term value system, but they basically have similar meanings. In general, a motivational system refers to an internal system of a human being that maintains the internal states or conditions that activate and influence voluntary behaviors of the human being. It is believed to enable the development of cognitive mechanisms to solve problems. For example, a child's hunger enables the child to get attention from its caregiver or to observe and learn the behavior of others who have successfully satisfied their hunger. The motivational system develops throughout the life experience. For example, an adult may choose to continue to work while he is hungry. Social morality, cultural tradition, and religious beliefs also play an important role in the motivational system of a human adult. Motivational systems are being studied by researchers in both animal (including human) learning and machine learning areas. While animal learning researchers are more interested in how motivational mechanisms are organized and developing in nature, machine learning researchers are, in addition, exploring how to build or develop motivational mechanisms in artificial agents.

- Supplied by Yilu Zhang and Juyang Weng

Editor: Yilu Zhang **Editorial assistant: Kim Thompson** **Logo: Joe Zimmerman**