The Animat

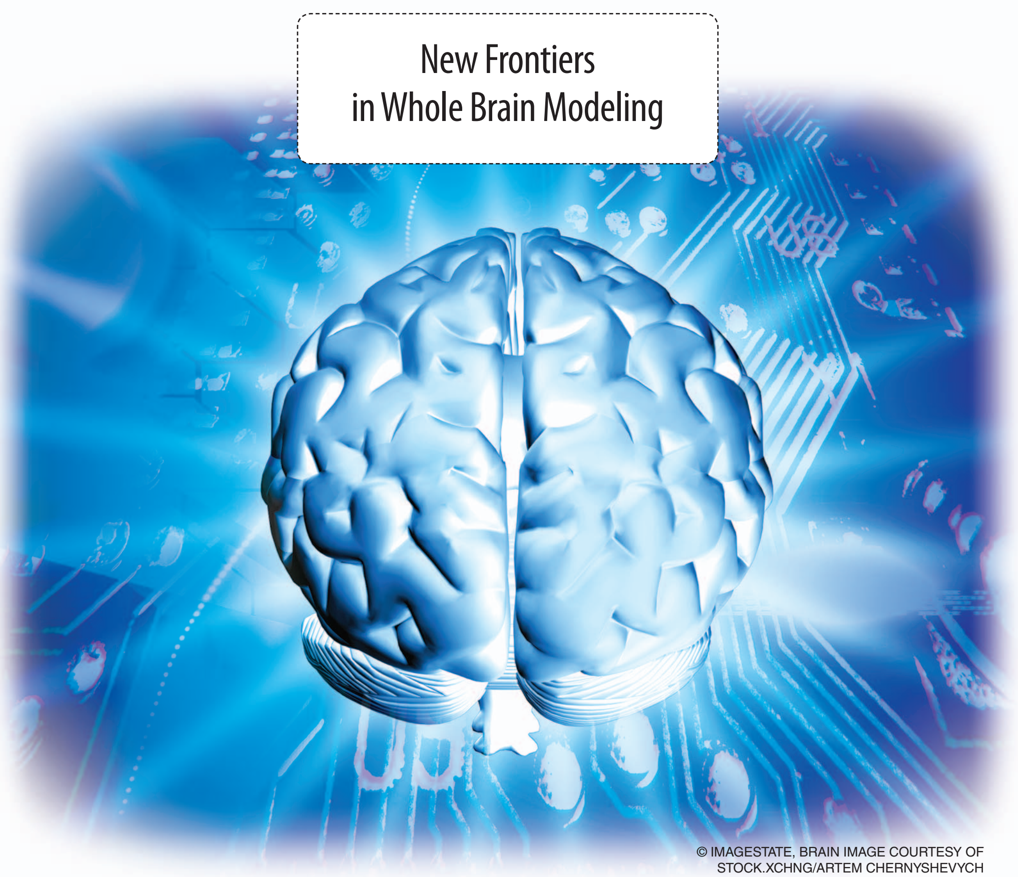

By Heather Ames, Ennio Mingolla, Aisha Sohail, Benjamin Chandler, Anatoli Gorchetchnikov, Jasmin Léveillé, Gennady Livitz, and Massimiliano Versace

The researchers at Boston University (BU)'s Neuromorphics Laboratory, part of the National Science Foundation (NSF)-sponsored Center of Excellence for Learning in Education, Science, and Technology (CELEST), are working in collaboration with the engineers and scientists at Hewlett-Packard (HP) to implement neural models of intelligent processes for the next generation of dense, low-power, computer hardware that will use memristive technology to bring data closer to the processor where computation occurs. The HP and BU teams are jointly designing an optimal infrastructure, simulation, and software platform to build an artificial brain. The resulting Cog Ex Machina (Cog) software platform has been successfully used to implement a large-scale, multicomponent brain system that is able to simulate some key rat behavioral results in a virtual environment and has been applied to control robotic platforms as they learn to interact with their environment.

MoNETA uses Cog to control an animat, which can be virtual simulations or physical robots.

Throughout the last century, people have speculated about whether machines would be able to think, learn, or behave in a humanlike way. Attempts to build thinking machines have had limited success to date. Many machines appear to behave like humans on specific tasks but completely fail when faced with

Digital Object Identifier 10.1109/MPUL.2011.2175638
Date of publication: 6 February 2012

The HP and BU teams are jointly designing an optimal infrastructure, simulation, and software platform to build an artificial brain.

unexpected situations because they are unable to learn when they encounter new environments, stimuli, or tasks. This situation is confounded by the problem that no one really knows how to build an artificial brain because we are not fully aware how our own brains work.

In the past few years, however, there have been advances in several fields of research that may lead to the building of an intelligent electronic brain. Neuroscience researchers have observed learning at the level of a single synapse and have gained considerable insight into how learning occurs at this scale. Psychologists have deeply analyzed the learning behaviors of humans in a wide variety of tasks and have also used other mammals, such as rats, to help understand intelligent behaviors. Computational neuroscientists have increasingly developed sophisticated neural models to describe both the underlying neural mechanisms and their resulting behaviors. Engineers have been hard at work developing a new computer hardware that will allow many more computations to be performed in a shorter amount of time while consuming less power.

One key hardware development is the advancement in understanding memristive devices. Memristive devices are two-terminal, nonlinear, passive electrical components that change their resistance as a function of voltage history. The technology offers potential benefits of much higher memory density, lower power consumption, and compatibility with the existing chip-manufacturing processes. Eventually, memristive devices may find their use as direct analogs to biological synapses in neuromorphic hardware even as a memory technology in a more conventional digital processor, but memristive technology will dramatically increase the performance and power efficiency of the processor.

These devices have allowed for a renewed hope in the artificial brain race because they show promise in changing the architecture of computer circuitry. Memristive materials are compatible with complementary metal–oxide–semiconductor (CMOS) processes, the standard technology for integrated circuits, thus allowing the designers to couple dense, memristive memories with conventional and widely available chips. This is truly an enabling technology for building artificial brains because the close proximity of memory and computation is a key reason as to why biological computation is so efficient.

In the past two years, the BU team has developed a novel modular approach in building the neural models that power an artificial brain. The macrostructure of the brain is initially specified with the goal of being able to swap in more refined neural circuits when they become available. Modules are currently being developed with increasing complexity and plugged into the architecture such that the artificial brain system will increase in its functional capabilities as development progresses. This whole brain system is called *modular neural exploring traveling agent* [MoNETA, Figure 1(c); [2]], and it runs on a special operating system called *Cog* [3], which is designed to provide a flexible, digital platform for the brain models and neuromorphic hardware to interact in a real or virtual environment. MoNETA uses Cog to control an animat, a term defining artificial animals, which can be virtual simulations or physical robots. Cog offers a way for the researchers to continue building models quickly at a relatively low cost while the development of the underlying technology advances.

The currently used MoNETA visual system involves parallel what and where processing streams to simultaneously identify and localize objects, respectively. An example of how MoNETA does this is the use of a dynamic attentional window in the where system that helps the what system focus on a restricted candidate area upon which to deploy limited attentional resources. Simultaneously, the what system, once an object is beginning to be identified, biases the where system to explore areas near the locus of attention that led to the object classification.

MoNETA has the ability to learn and have higher-level cognitive skills. To achieve this goal, MoNETA has a simulated motivational system that is modeled after animal drives to evaluate the utility of objects in its surroundings. MoNETA implements reinforcement learning models for motivation and reward, which in turn bias how the animat selects goals in the environment. Examples of biological drives are lack of comfort and curiosity toward unexplored spatial locations. At any given time, the simulated equivalents of these biological drives compete to control the motor system and move the animat toward a desired location. In addition, MoNETA learns the physical locations of rewards and punishments via reinforcement learning using the feedback from the environment. MoNETA then uses the current

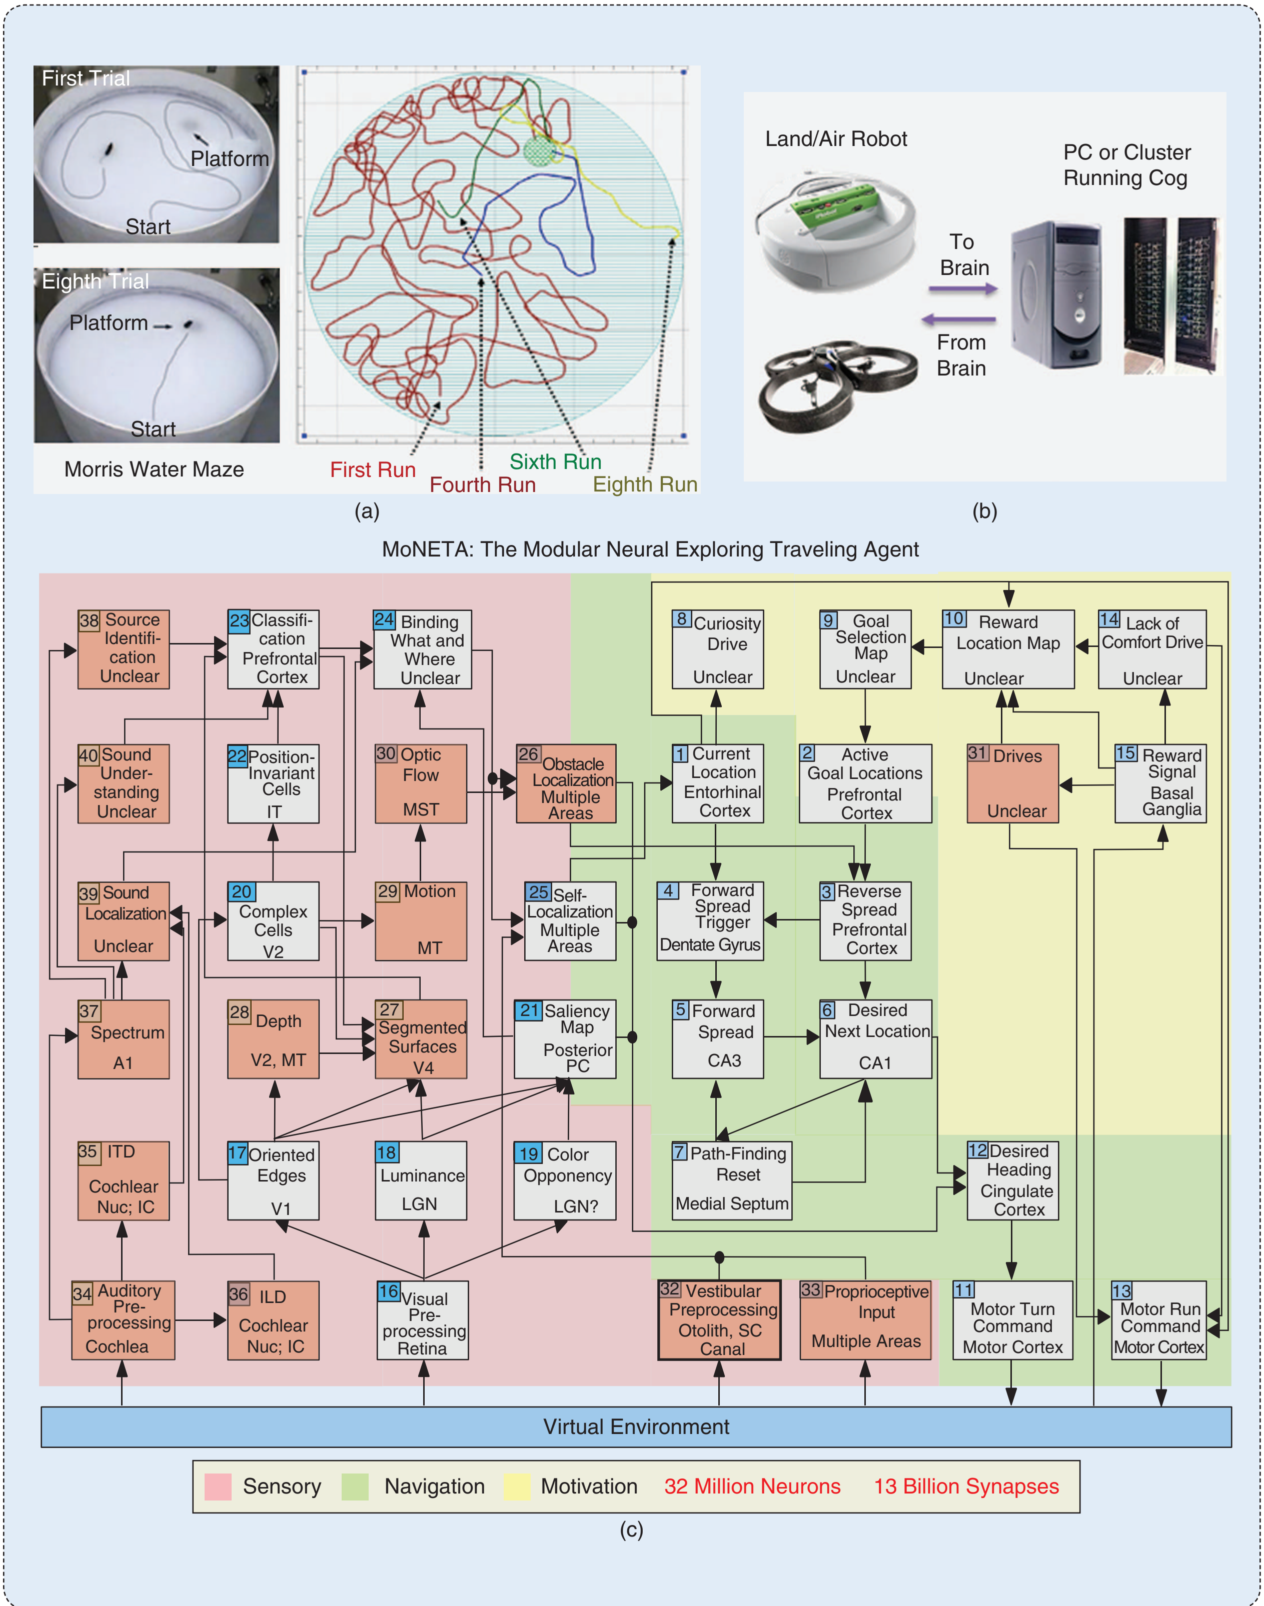

FIGURE 1 (a) MoNETA learns to perform the Morris water maze task, a typical rat behavioral paradigm where a rat learns the position of a submerged platform. In the first trial, MoNETA explores its environment driven by simulated motivational states, or drives, representing the biological equivalents of lack of comfort and curiosity. Once it, by chance, swims on top of the submerged platform (green), MoNETA learns its position by using some visual landmarks at the border of the pool. As training progresses, MoNETA is able to directly swim from its current position to the platform. (b) MoNETA can be run on robotic platforms, both land and air, by wirelessly interfacing the robot to a workstation of a cluster where Cog runs the MoNETA brain. (c) The system diagram of MoNETA. It is a whole-brain-system model consisting of sensory, motivation, and navigation areas interfaced with a virtual environment or robotic platform. The first version of MoNETA consisted of approximately 32 million neurons and 13 billion synapses.

There have been advances in several fields of research that may lead to the building of an intelligent electronic brain.

location and learned goal location to approach the goal via a bidirectional graph search that determines the optimal path to the desired location even in cluttered environments.

MoNETA also integrates sensory information (currently, vision, touch, and proprioception) into higher-order representations of its emerging reality and is able to react to novel situations not explicitly programmed within the software.

MoNETA v1.0 is completed and has been tested in an animat negotiating a virtual Morris water maze task [Figure 1(a)]. The Morris water maze [1] is a task used to probe the navigation skills of a rodent. In this classic task, the rat is placed in a water tank and has to use visual cues to locate a submerged platform and swim toward it. The rat is motivated to find the platform because it does not like being wet. Researchers have studied this task in detail, so we know a great deal about the brain areas that a rat uses when completing this task. Although it is an apparently simple task, solving the water maze requires the integrated simulation of brain areas subserving object recognition and localization, touch, proprioception, goal selection, motivation, and navigation.

The development of new hardware technologies is still in the early phase of research and testing, so MoNETA is powered by simulated architectures that take advantage of a heterogeneous arrangement of computer processors. Cog provides the glue between the hardware and the neural models, which in turn allows for a seamless integration of new hardware as it is developed.

The success of the animat on the Morris water maze task marks the end of the first phase of the project. Initial testing of MoNETA was also done on robotic platforms in a real environment [Figure 1(b)]. The next phase will involve running the animat in several other experimental rodent-maze paradigms and expanding the complexity of the neural models powering the animat. In parallel, advances in software and hardware architectures brought forward by our colleagues at HP will bring us closer to being able to implement these models in robotic and mobile platforms.

Although MoNETA is far from a thinking learning machine scalable to human intelligence, it does represent an important first milestone in the development of such an artificial brain. It models biological brain functions, replicates rodent behavior in a simple paradigm, and takes into account new and vibrant advances in hardware that have come closer to mimicking biology than ever before.

Acknowledgments

The authors thank Cynthia Bradford for her help with manuscript preparation. The work was supported in part by the CELEST, an NSF Science of Learning Center (NSF SBE-0354378 and NSF OMA-0835976). This work was also partially funded by the Defense Advanced Research Projects Agency (DARPA) Systems of Neuromorphic Adaptive Plastic Scalable (SyNAPSE) program, contract HR0011-09-3-0001. The views, opinions, and/or findings contained in this article are those of the authors and should not be interpreted as representing the official views or policies, either expressed or implied, of the DARPA, the Department of Defense, or the NSF.

Heather Ames (heather.m.ames@gmail.com), Aisha Sohail (aishas110@gmail.com), Benjamin Chandler (bchandle@gmail.com), *and* ***Gennady Livitz (glivitz@gmail.com)*** *are with the Center of Excellence for Learning in Education, Science, and Technology (CELEST), Boston University, Massachussetts.* ***Ennio Mingolla (ennio@cns.bu.edu), Anatoli Gorchetchnikov (anatoli@cns.bu.edu), Jasmin Léveillé (jasminl@cns.bu.edu),*** *and* ***Massimiliano Versace (versace@cns.bu.edu)*** *are with Neuromorphics Lab, Boston University, Boston, Massachussetts.*

References

[1] R. G. M. Morris, "Spatial localization does not depend on the presence of local cues," *Learn. Motivation*, vol. 12, no. 2, pp. 239–260, 1981.

[2] M. Versace and B. Chandler, "MoNETA: A mind made from memristors," *IEEE Spectr.*, vol. 47, no. 12, pp. 30–37, Dec. 2010.

[3] G. Snider, R. Amerson, D. Carter, H. Abdalla, S. Qureshi, J. Laveille, M. Versace, H. Ames, S. Patrick, B. Chandler, A. Gorchetchnikov, and E. Mingolla, "Adaptive computation with memristive memory," *IEEE Computer*, vol. 44, no. 2, pp. 21–28, 2011.