Learning in and from Brain-Based Devices
Gerald M. Edelman, *et al.*
Science **318**, 1103 (2007);
DOI: 10.1126/science.1148677

PERSPECTIVE

Learning in and from Brain-Based Devices

Gerald M. Edelman

Biologically based mobile devices have been constructed that differ from robots based on artificial intelligence. These brain-based devices (BBDs) contain simulated brains that autonomously categorize signals from the environment without a priori instruction. Two such BBDs, Darwin VII and Darwin X, are described here. Darwin VII recognizes objects and links categories to behavior through instrumental conditioning. Darwin X puts together the "what," "when," and "where" from cues in the environment into an episodic memory that allows it to find a desired target. Although these BBDs are designed to provide insights into how the brain works, their principles may find uses in building hybrid machines. These machines would combine the learning ability of BBDs with explicitly programmed control systems.

In the last several decades, great progress has been made in describing how the separate components of the nervous system function. Less progress has been made in obtaining a global picture of higher brain functions such as learning and memory. This picture must take into account that the brain is embodied and that the body and brain act together in a real-world environment. Here, I briefly describe a synthetic approach to elucidating how integrated nervous systems operate by constructing brain-based devices (BBDs). These devices differ fundamentally from robots that are controlled by built-in artificial intelligence consisting of explicit programmed instructions (*1*). Instead, like real animals, BBDs must learn autonomously to categorize signals from the environment without a priori instruction.

BBDs are constructed according to a procedure called synthetic neural modeling (*2*). In such an approach a detailed brain is simulated in a computer and controls a mobile platform containing a variety of sensors and motor elements. In modeling the properties of real brains, efforts are made to simulate vertebrate neuronal components, neuroanatomy, and dynamics in detail (Fig. 1A). As a result of the exploration of a real world environment by a BBD, the BBD develops adaptive behavior through processes mimicking those underlying learning in animals. In its close mimicry of vertebrate neural systems, the BBD approach stands in contrast to other more functionally oriented neurorobotic models (*3*).

To bring out the differences from programmed robots, I shall describe the composition and behavior of two BBDs called Darwin VII and Darwin X. These devices are named after the great biologist Charles Darwin to emphasize the fact that their brains learn by selection from a repertoire of many different simulated neural circuits and do not depend on explicitly programmed instructions.

Darwin VII (Fig. 1A) has a mobile base fitted with a video camera for vision, a pair of microphones for audition (contained in two plastic cups), and a gripper device that can grab steel blocks in its environment, each painted with either stripes or blobs. The BBD's brain consists of a visual system, an auditory system, a "taste" system in the gripper (which measures the conductance of a gripped block), a motor system capable of triggering movement, and a value system.

The value system in the BBD's brain signals the salience of environmental cues and leads to rewarding or aversive responses that enable the device's learning behavior to be adaptive. The simulated neurons underlying these systems are linked in circuits modeled on known anatomy. The connections within these neural circuits (synapses) can change in their strength after receiving sensory signals. The patterns of these changes are unique to each individual BBD because they reflect that device's past behavior. Darwin VII's complete nervous system contains about 20,000 simulated neuronal units linked together by 450,000 synaptic connections [see supporting online material (SOM)]. As the BBD was exploring the environment, the complex responses of these units to signals from the environment were modeled in a large computer cluster and were radioed to Darwin VII to direct motor activity. Consequent changes in the sensory inputs to the device were sent back to the simulated brain in a dynamic fashion, allowing smooth movement in real time (*4*).

The basic modes of behavior of Darwin VII consist of visual exploration and tracking, gripping and tasting, and two innate reflex responses: appetitive and aversive. Figure 1A shows Darwin VII approaching a steel block with stripes (detected by the camera) that was arbitrarily constructed to have good taste (high conductance). Gripping the block sent appetitive signals to the value system in the brain. Blocks containing blobs had low conductance (bad taste) and sent aversive signals to that system.

In a series of conditioning experiments, Darwin VII autonomously explored its environment and learned to associate the taste value of the blocks with their visual patterns. Appetitive and aversive responses were initially triggered by taste, but after about 10 encounters with a population containing both types of block, the BBD's responses were triggered by vision alone. More than 90% of the time, after training, the BBD continued to grip and taste striped blocks but actually backed away from blocks with blobs (Fig. 1B).

This autonomous development of patterned behavior by instrumental or operant conditioning requires sophisticated visual anatomy (see SOM for details). The inset of Fig. 1A shows the patterns of activity of four simulated areas as the BBD moves. The upper left square illustrates a pattern of activity in an early visual brain area that responds in a striped pattern to a striped block. The upper right square shows a firing pattern of neuronal units in an integrative higher visual region corresponding to the inferotemporal area of the brain, an area in animal brains that responds differentially to different object categories. This particular pattern is dependent on the BBD's history of encounters with striped blocks and thus is unique to that history (an identical Darwin VII with a different history would have its own characteristic but equally correlative pattern). The colored areas in the lower right-hand square signal positive value responses as appetitive taste, even before gripping. Appetitive responses in brain circuits prompt motor action, resulting in approach and grabbing of the block. The lower left-hand square would signal aversive responses, but none were experienced in this sequence. If Darwin VII encountered the block with a blob pattern (see block at bottom of main panel), aversive activity would be signaled in that square.

A critical feature of such a BBD is that, after a period of training and behavior, the activity of all neuronal units, synapses, and circuits can be recorded and examined in detail. This cannot be achieved in experiments on living animals. As I shall mention later, the patterns obtained from this type of exhaustive analysis can be of great value in analyzing brain function. To extend such analyses, experiments can be performed to examine experience-dependent changes in the BBD's perceptual categorization and learning. These changes have included learning followed by reversal learning; i.e., learning after switching appetitive and aversive correlations.

Darwin VII's memory after learning did not reflect sequences of events. Given this limitation, my colleagues and I asked whether we could model long-term episodic memory in a BBD. This

The Neurosciences Institute, 10640 John Jay Hopkins Drive, San Diego, CA 92121, USA. E-mail: edelman@nsi.edu

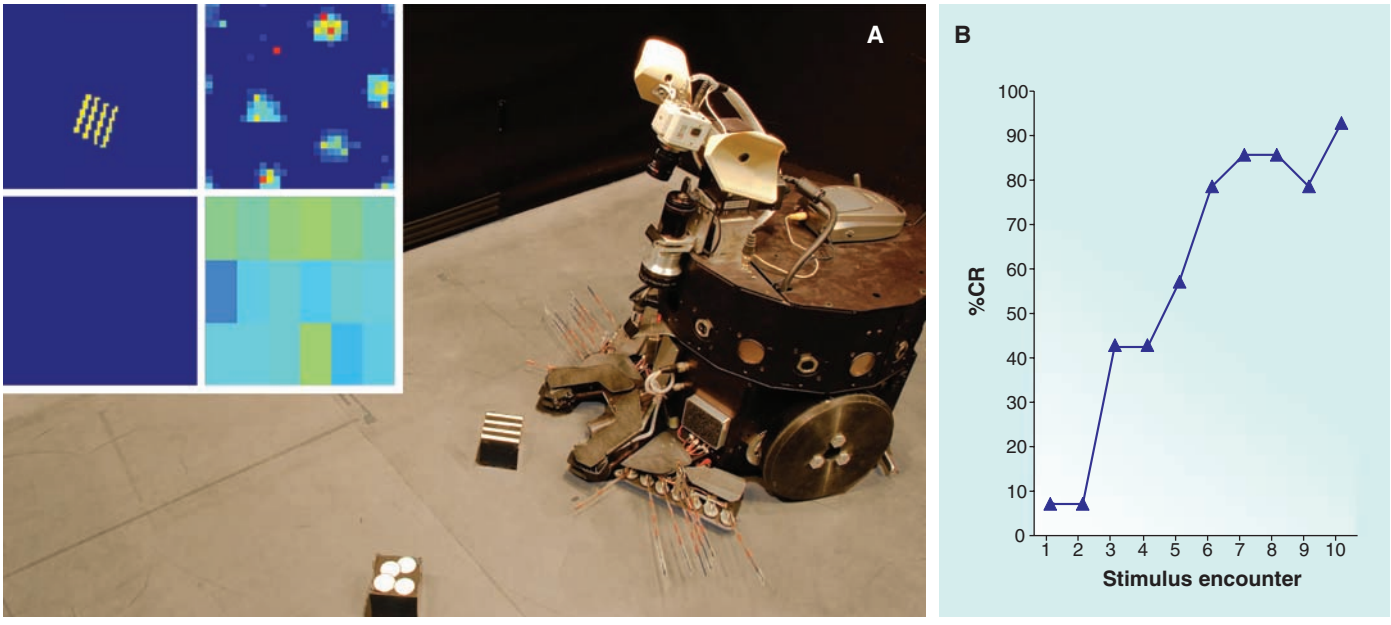

Fig. 1. (**A**) Darwin VII approaching a striped block after training. The mobile device is equipped with a charge-coupled device camera for vision, microphones for hearing, and a gripper capable of sensing levels of conductance for taste. (Inset) Patterns of activity of four simulated areas as the BBD moves. Upper left, activity in an early visual brain area responding in a striped pattern to a striped block; upper right, firing pattern of neuronal units in an integrative higher visual region corresponding to the inferotemporal area of the brain; lower right, colored areas signal positive value responses as appetitive taste, even before gripping; lower left, aversive responses (none in this sequence). See text for further details. (**B**) Percentage of correct responses (CRs) as a function of stimulus block encounters during learning trials for visual conditioning.

memory of sequences of events in real time is known to depend critically on a brain structure called the hippocampus. Patients who have lesions of the hippocampus bilaterally can remember episodic events in their early lives, but after the occurrence of the lesion can no longer convert short-term memory into long-term episodic memory. In rats, the hippocampus is also known to be essential for the successful integration of cues allowing navigation to a goal in a remembered environment. One test of this capability is provided by the so-called Morris water maze. A rat is placed in a tank of milky water that has a hidden platform below the surface. The animal will traverse various paths until it locates the platform, where it will then rest. Having seen various cues on the wall of its surroundings during its traverse, the rat possesses episodic recall and can subsequently locate the hidden platform directly upon being placed again in any part of the tank. Lesions of the hippocampus compromise this ability.

We decided to model the hippocampus in the brain of a BBD called Darwin X (*5*). This device was then tested in a dry version of the maze (Fig. 2A). An enclosure was constructed with a black floor and walls. Each wall had differently colored paper strips of varying widths that could act as visual cues. At one location we placed a "hidden" platform. The visual system of Darwin X could not detect the platform, but the BBD was equipped with an infrared (IR) detector that would respond only when the device was directly over the platform. That detector's signal then triggered a positive value response. This response is analogous to that of a rat in a water maze sensing a solid platform under its feet.

Darwin X was given the task of finding the hidden platform, and after 8 to 10 traverses during which it detected different cues on the

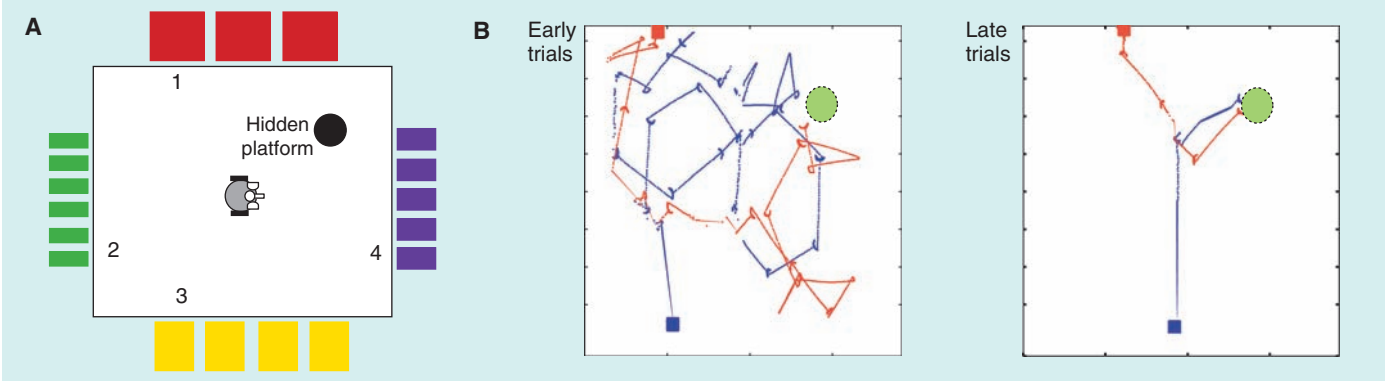

Fig. 2. (**A**) Schematic layout of the enclosure used for the hidden-platform task. The enclosure measured 16 by 14 feet, with black walls and flooring. Pieces of differently colored paper of varying widths (to act as different cues) were hung on each of the walls. A hidden circular platform—24 inches in diameter and made of reflective black paper—was placed in the center of a quadrant in the enclosure. Each trial began in one of four starting locations (numbers 1 to 4 in the diagram). (**B**) Behavioral performance in the hidden-platform task indicated by trajectories of a subject during the training paradigm. Green circles denote the location of the platform during the training trials, red and blue squares denote the starting locations, and red and blue lines indicate trajectories during individual early and late trials. The late trials showed more or less direct paths, regardless of the starting point. [Adapted from (*4*)]

wall, it would directly go to the platform from any starting point (Fig. 2B). Indeed, if the hidden platform was removed, a trained Darwin X would, like a trained rat, concentrate its searches in the area previously occupied by the platform.

An analysis of the neural responses of the simulated hippocampus revealed that they closely resembled those in the living animal. The simulated brain of Darwin X had 50 different neural areas, 90,000 neuronal units, and 1.4 million synaptic connections. We could record the response of every neuronal unit and connection in the BBD, a procedure not possible in a living animal. After training, we could pick any neuronal unit that sent signals to the motor system and trace back the types and strengths of connection of all neuronal units that caused the firing activity of this reference unit. An examination of the resultant backtrace network revealed a very large number of different possible paths and circuits leading to the firing of a single chosen reference unit (*5*). Thus, the system showed degeneracy, the property according to which different structures can lead to the same response or output (*6*). This observation suggests that degeneracy might also be a common property of the neural networks of living animals. By providing these insights, synthetic neural modeling has enhanced our efforts to understand how the human brain works.

In addition to these fundamental issues, there is a practical implication for the field of robotics of the work on BBDs. It is now possible to construct hybrid machines incorporating the principles of BBDs together with engineering principles that rely on programmed instructions. An example is a robotic soccer-playing device that we constructed on the platform of a Segway Human Transporter at the suggestion of the Defense Advanced Research Program Agency. This device was designed to play together with a human teammate in the 2005 U.S. Open Robocup Championship in Atlanta. It had a video camera for recognizing objects on the field (balls, teammates, goals, etc.), IR sensors and a laser range finder to detect the ball, and a ball-capturing and -kicking device. The Segway platform's behavior was guided by a neural simulation in an on-board computer that received inputs from the various sensors and generated motor signals to the Segway's wheels. In addition to the simulated brain containing neuronal units, it had a set of programs (like a conventional robot) guiding some of its reflex responses. Playing against an excellent team from Carnegie Mellon University that used a Segway device based on artificial intelligence, our team won all five games (*7*). We attribute this largely to the ability of our Segway device to learn by experience before the game.

We expect that, although BBDs were initially designed to further our understanding of the human brain, their principles may complement those currently used in various engineering approaches. Clearly, there are several other areas of robotics that, to some degree, have taken account of biological principles or that may contribute to the effectiveness of hybrid designs. This prompts a detailed comparison of our work with two growing fields of robotics. The first, evolutionary robotics (*8*), views robots as autonomous artificial organisms that develop skills by selection after interacting with their environment. The second, probabilistic robotics (*9*), is concerned with perception and control by robots in the face of uncertainty.

BBD design is still in its early stages. It will be greatly enhanced by the development of small, very powerful computers that are capable of direct placement on board the platform of each device. A far-off goal of BBD design is the development of a conscious artifact (*10*). Although machine consciousness (*11*) is at best a remote prospect, the fact that we can build BBDs that are capable of perceptual categorization with sophisticated memory systems provides an initial basis for what a decade ago would have been considered science fiction.

References

1. S. Thrun, G. Sukhatme, S. Schaal, O. Brock, Eds., *Robotics: Science and Systems I* (MIT Press, Cambridge, MA, 2005).
2. G. M. Edelman *et al*., *Proc. Natl. Acad. Sci. U.S.A.* **89**, 7267 (1992).
3. A. K. Seth, O. Sporns, J. L. Krichmar, *Neuroinformatics* **3**, 167 2005.
4. J. L. Krichmar, G. M. Edelman, *Artif. Life* **11**, 63 (2005).
5. J. L. Krichmar, D. A. Nitz, J. A. Gally, G. M. Edelman, *Proc. Natl. Acad. Sci. U.S.A.* **102**, 2111 (2005).
6. G. M. Edelman, J. A. Gally, *Proc. Natl. Acad. Sci. U.S.A.* **98**, 13763 (2001).
7. See www.nsi.edu/nomad/segway and www.cs.cmu.edu/~robosoccer/segway for videos of the games.
8. S. Nolfe, D. Floreano, *Evolutionary Robotics: The Biology, Intelligence, and Technology of Self Organizing Machines* (MIT Press, Cambridge, MA, 2000).
9. S. Thrun, W. Burgard, D. Fox, *Probabilistic Robotics* (MIT Press, Cambridge, MA, 2005).
10. G. M. Edelman, *Proc. Natl. Acad. Sci. U.S.A.* **100**, 5520 (2003).
11. O. Holland, Ed., *Machine Consciousness* (Imprint Academic, Exeter, UK, 2003).

Supporting Online Material

www.sciencemag.org/cgi/content/full/318/5853/1103/DC1
SOM Text
Figs. S1 to S4
References

10.1126/science.1148677