

What Can AI Get from Neuroscience?

Steve M. Potter

Laboratory for Neuroengineering
Department of Biomedical Engineering
Georgia Institute of Technology
313 Ferst Dr NW, Atlanta, GA, USA 30332-0535
steve.potter@bme.gatech.edu
<http://neuro.gatech.edu>

Abstract. The human brain is the best example of intelligence known, with unsurpassed ability for complex, real-time interaction with a dynamic world. AI researchers trying to imitate its remarkable functionality will benefit by learning more about neuroscience, and the differences between Natural and Artificial Intelligence. Steps that will allow AI researchers to pursue a more brain-inspired approach to AI are presented. A new approach that bridges AI and neuroscience is described, Embodied Cultured Networks. Hybrids of living neural tissue and robots, called hybrot, allow detailed investigation of neural network mechanisms that may inform future AI. The field of neuroscience will also benefit tremendously from advances in AI, to deal with their massive knowledge bases and help understand Natural Intelligence.

Keywords: Neurobiology, circular causality, embodied cultured networks, animats, multi-electrode arrays, neuromorphic, closed-loop processing, Ramon y Cajal, hybrot.

1 Introduction

An alien power plant was unearthed in a remote South American jungle. After excavating and dusting it off, the archeologists flip the switch, and it still works! It generates electricity continuously without needing fuel. Wouldn't we want to make more of these power plants? Wouldn't we want to know how this one works? What if the scientists and engineers who design power plants saw photos of the locals using electricity from the alien power plant, and knew it reliably powers their village. Yet they ignore this amazing artifact, and feel it has little relevance to their job. Although this scenario seems implausible, it is analogous to the field of AI today. We have, between our ears, a supremely versatile, efficient, capable, robust and intelligent machine that consumes less than 100 Watts of power. If AI were to become less artificial, more brain-like, it might come closer to accomplishing the feats routinely carried out by *Natural Intelligence* (NI). Imagine an AI that was as adept as humans at speech and text understanding, or reading someone's mood in an instant. Imagine an AI with human-level creativity and problem solving. Imagine a dexterous AI, which could precisely and adaptively manipulate or control physical artifacts such as violins, cars, and balls. Humans, thanks to our complex nervous system, are especially

good at interacting with the world in real time in non-ideal situations. Yet, little attention in the AI field has been directed toward actual brains. Although many of the brain's operating principles are still mysterious, thousands of neuroscientists are working hard to figure them out.¹

Unfortunately, the way neuroscientists conduct their research is often very reductionistic [1], building understanding from the bottom up by small increments. A consequence of this fact is that trying to learn, or even keep up with, neuroscience is like trying to drink from a fire hose. General principles that could be applied to AI are hard to find within the overwhelming neuroscience literature.

AI researchers, young and old, might do well to become at least somewhat bilingual. Taking a neuroscience course or reading a neuroscience textbook would be a good start. Excellent textbooks include (among others) *Neuroscience* [2], *Neuroscience: Exploring the Brain* [3], and *Principles of Neural Science* [4]. There are several magazines and journals that allow the hesitant to gradually immerse themselves into neuroscience, one toe at a time. These specialize in conveying general principles or integrating different topics in neuroscience. In approximate order of increasing difficulty, some good ones are: *Discover*, *Science News*, *Scientific American Mind*, *Cerebrum*, *Behavioral and Brain Sciences (BBS)*, *Trends in Neuroscience*, *Nature Reviews-Neuroscience*, and *Annual Review of Neuroscience*. BBS deserves special mention, because of its unusual format: a 'target article' is written by some luminary, usually about a fairly psychological or philosophical aspect of brains. This is followed by in-depth commentaries and criticisms solicited from a dozen or more other respected thinkers about thinking. These responses provide every side of a complex issue, and often include many of the biological foundations of the cognitive functions being discussed. The responses are followed by a counter-response from the author of the target article. BBS is probably the best scholarly journal that regularly includes and combines contributions from both neuroscientists and AI researchers.²

In this networked era, the internet can be a cornucopia, or sometimes, a Pandora's Box for AI researchers who want to learn about real brains. Be wary of web pages expounding brain factoids, unless there is some form of peer review that helps maintain the quality and integrity of the information. Wikipedia is rapidly becoming an extremely helpful tool for getting an introduction to any arcane topic, and has an especially elaborate portal to Neuroscience.³ Caution: it is not always easy to find the source or reliability of information given there. A more authoritative source on the fields of computational neuroscience and intelligence is *Scholarpedia*.⁴ The Society

¹ I will define neuroscience as all scientific subfields that aim to study the nervous system (brain, spinal cord, and nerves), including neurophysiology, neuropathology, neuropharmacology, neuroendocrinology, neurology, systems neuroscience, neural computation, neuroanatomy, neural development, and the study of nervous system functions, such as learning, memory, perception, motor control, attention, and many others. Neurobiology is thought of today as the basis of all neuroscience (ignoring some lingering dualism) and the terms are often used interchangeably.

² BBS Online: <http://journals.cambridge.org/action/displayJournal?jid=BBS>

³ <http://en.wikipedia.org/wiki/WP:NEURO>

⁴ <http://www.scholarpedia.org>

for Neuroscience (SFN) website⁵ is an excellent and reliable source of introductory articles about many neuroscience topics. The SFN consists of over 30,000 (mostly American) neuroscientists who meet annually and present their latest research to each other. All of the thousands of abstracts for meetings back to the year 2000 are searchable on the Annual Meeting pull-down. Although not itself a repository of introductory neuroscience material, the Federation of European Neuroscience Societies website⁶ is a good jumping-off point for all things Euro-Neuro.

2 What Do We Already Know About NI (Natural Intelligence) That Can Inform AI?

2.1 Brains Are Not Digital Computers

John von Neumann, the father of the architecture of modern digital computers, made a number of thought-provoking and influential analogies in his book, "The Computer and the Brain." [5] The brain-as-digital-computer metaphor has proven quite popular, and often gets carried too far. For example, a neuron's action potential⁷ is often referred to by the AI field as a biological implementation of a binary coding scheme. This and other misinterpretations of brain biology need to be purged from our thoughts about how intelligence may be implemented. Even with our rudimentary conception of how it is implemented in brains, there are clear differences between computers and brains, such as:

2.2 Brains Don't Have a CPU

The brain's processor is neither "central" nor a "unit". Its processing capabilities seem to be distributed across the entire volume of the brain. Some localized regions specialize in certain types of processing, but not without substantial interaction with other brain areas [6].

2.3 Memory Mechanisms in the Brain Are Not Physically Separable from Processing Mechanisms

Recent research has shown that in recalling a memory, similar brain regions are activated as during perceiving [7]. This may be because an important part of perceiving is comparing sensory inputs to remembered concepts. Memories are dynamic, and continually re-shaped by the recall process [8]. A computer architecture that unites the processor, RAM, and hard disk into one and the same substrate might be far more efficient. An architecture that implements memory as a dynamic process rather than a static thing may be more capable of interacting in real time with a dynamic world.

⁵ <http://www.sfn.org>

⁶ <http://fens.mdc-berlin.de>

⁷ Action potentials are regenerative electrical impulses that neurons evolved to send information across long axons. They involve a fluctuation of the neuron membrane potential of ~0.1 V across a few milliseconds.

2.4 The Brain Is Asynchronous and Continuous

The computer is a rare type of artifact that has well-defined (discontinuous) states [9], thanks to the fact that its computational units are always driven to their binary extremes each tick of the system clock. There are many brain circuits that exhibit oscillations [10], but none keeps the whole brain in lock-step the way a system clock does for a digital computer. The phase of some neural events in relation to a circuit's ongoing oscillation is used to code for specific information [11], and phase is a continuous quantity.

2.5 With NI, the Details of the Substrate Matter

Digital computers have been very carefully designed so that the details of their implementation don't influence their computations. Vacuum tubes, discrete transistors, and VLSI transistors, since they all speak Boolean, can all run the same program and produce the same result. There is a clear, intentional separation between the hardware and the software. All neuroscience research so far suggests this separation does not exist in the brain.

How do the details of its substrate influence the brain's computations? Every molecule that makes up the brain is in continuous motion, as with all liquids. The lipid bilayer that comprises the neuron's membrane is often referred to as a 2-dimensional liquid and is part of the neural wetware. The detailed structure of the proteins that make up brain cells can only be determined when they are crystallized in a test tube, that is, purified and stacked into unnatural, static, repeating structures that form good x-ray diffraction patterns. In their functional form, proteins (and all brain molecules) are jostling around, continuously bombarded by the cytoplasm or cerebrospinal fluid that surrounds them, like children frolicking in a pen full of plastic balls. Small details about neurons' structure, such as the morphing of tiny (micron-sized) synaptic components called dendritic spines [12], or the opening and closing of voltage-sensitive or neurotransmitter-sensitive ion channels, affect their function at every moment. All that movement of molecules and parts of cells is the substrate of NI, facilitating or impeding communication between pairs of brain cells and across functional brain circuits. Why should AI researchers concern themselves with the detailed, molecular aspects of brain function? Because, fully duplicating brain functionality may only be possible using a substrate as complex and continuous as living brain cells and their components are.

That disappointing possibility should not keep us from trying at least to duplicate *some* brain functionality by taking cues from NI. Carver Mead, Rodney Douglas, and other neuromorphic engineers have designed useful analog circuits out of CMOS components that take advantage of more of the physics of doped silicon than just its ability to switch from conducting to non-conducting states [13]. The continuous "inter-spike interval" between action potentials in neurons is believed to encode neural information [14] and also seems to be responsible for some of the brain's learning abilities [15]. Neuromorphic circuits that use this continuous-time pulse-coding scheme [13, 47] may be able to process sensory information faster and more efficiently than could digital circuits.

2.6 NI Thrives on Feedback and Circular Causality

The nervous system is full of feedback at all levels, including the body and the environment in which it lives; it benefits in a quantifiable way from being embodied and situated [16, 17]. Unlike many AI systems, NI is highly interactive with the world. Human-engineered systems are more tractable when they employ assembly-line processing of information, i.e., to take in sense data, then process it, then execute commands or produce a solution. Most sensory input to living systems is a dynamic function of recent or ongoing movement commands, such as directing gaze, walking, or reaching to grasp something. With NI, this active perception and feedback is the norm [17, 18]. Animal behaviors abound with circular causality, new sensory input continuously modulating the behavior, and behavior determining what is sensed [19]. One beautiful example of active perception that humans are especially good at is asking questions. If we don't have enough information to complete a task, and a more knowledgeable person is available, we ask them questions. New AI that incorporates question-asking and active perception can solve problems quickly that would take too long to solve by brute force serial computation [16, 20].

There are few brain circuits that involve unidirectional flow of information from the sensors to the muscles. The vast majority of brain circuits make use of what Gerald Edelman calls reentry [21]. This term refers to complex feedback on many levels, which neuroscientists have only begun to map, let alone understand. Neuroscience research suggests that a better understanding of feedback systems with circular causality would help us design much more flexible, capable, and faster AI systems [9].

2.7 NI Uses LOTS of Sensors

One of the most stunning differences between animals and artificial intelligences is the huge number of sensors animals have. NI mixes different sensory modalities to enable rapid and robust real-time control. Our brains are very good at making the best use of whatever sense data are available. Without much training, blind people can deftly navigate unfamiliar places by paying attention to the echoes of sounds they make, even while mountain-biking off road!⁸[22] Bach y Rita's vibrotactile display placed a video camera's image onto a blind person's skin, in the form of a few hundred vibrating pixels. By actively aiming the camera, the user could "see" tactile images via their somatosensory system, allowing them to recognize faces and to avoid obstacles [23, 24]. The continuous flow of information into the brain from the sense organs is enormous. To make AI less artificial, we could strive to incorporate as much sensing power as we dare imagine. When AI adopts a design philosophy that embraces, rather than tries to minimize high-bandwidth input, it will be capable of increasingly more rapid and robust real-time control.

2.8 NI Uses LOTS of Cellular Diversity

There are more different types and morphologies of cells in the brain than in any other organ, perhaps than all organs and tissues combined. Many of these were catalogued by neuroscientist, Santiago Ramon y Cajal a century ago (Fig. 1) [25, 26], but more

⁸ <http://www.worldaccessfortheblind.org/>




Fig. 1. Neurons and circuits traced a century ago by Spanish neuroanatomist, Santiago Ramon y Cajal (pictured in center). This montage depicts only a few of the many types of neuron morphologies found throughout the nervous system. (Adapted with permission from Swanson & Swanson ©1990 MIT Press).

are still being discovered [27]. Another sign of the brain's complexity is the large amount of genetic information that allows it to develop and function. Both mice and men have ~30,000 genes in their genome, and over 80% of these are active in the

brain.⁹ All this cellular complexity and diversity may be crucial in creating an intelligent processor that is general-purpose and highly adaptable.

2.9 NI Uses LOTS of Parallelism

The brain's degree of parallelism is not rivaled by any human-made artifact. There are about 100 billion neurons in our brains, each connected to 1,000-10,000 others with over 200,000 km of axons, which we have barely begun to map [28]. The brain's circuits seem to have small-world connectivity [29], i.e., many local connections and relatively few, but crucial, long range connections. The latter integrate the activities of cooperating circuits running simultaneously and asynchronously. Although our knowledge of this elaborate connectivity is rudimentary, some general principles, such as small-world connectivity, could make future AI more capable.

2.10 Delays Are Part of the Computation

It is sometimes mistakenly stated that neurons are slow computational elements, since they fire action potentials at a few hundred Hz at most. The parallelism mentioned above is one way to enable rapid computation with "slow" elements. Modern computers, which are not very parallel if at all, reduce computation and transmission delays in every way possible, from shorter leads, to faster clocks, to backside caches. Any time spent getting information from here to there in a digital computer is viewed as a wasteful impediment to getting on with the computation of the moment. In the brain, delays are not a problem, but an important part of the computation. The subtle timing of action potentials carries information about the dynamics and statistics of the outside world [30]. The relative timing of arrival of two action potentials to the postsynaptic neuron determines whether the strength of their synapse is incrementally increased or decreased [31]. These pulse timings are analog quantities. The brain computes with timing, not Boolean logic [32]. Brain-inspired AI of the future will be massively parallel, have many sensors, and will make good use of delays and the dynamics of interactions between analog signals [33].

3 What Do We *Not* Know About How Brains Work, But Could Learn?

To realize this dream of AI that is closer to NI, there are a number of important questions about how brains work that must be pursued, such as: What is a memory? How do biological networks work? The neurons and glial cells both store and process information in a spatially distributed manner. But we have only a very vague and fuzzy idea of just how they do that. The Blue Brain Project is setting a giant supercomputer (the son of Deep Blue) to the task of simulating just one cortical minicolumn of a few thousand neurons [34]. There is a lot going on at the level of networks that we don't even have the vocabulary to think about yet. Neurobiologists all believe that memories are stored by changes in the physical structure of brain cells such as increases in the number of branches or spines on a neuron's dendritic tree. We

⁹ See the Allen Institute Brain Atlas, <http://www.brainatlas.org/aba/index.shtml>

don't all agree about what those changes might be, let alone how the changes are executed when salient sensory input is received. As hinted by Ramon y Cajal's drawings (Fig. 1), neurons have a stunning diversity of morphologies [35]. There is evidence that some aspects of their shape are altered by experience [36-38]. But how that relates to a memory being stored is not known.

4 New Neuroscience Tools

A new type of experimental animal, called the Hybrot, is taking shape in the Laboratory for Neuroengineering at the Georgia Institute of Technology. This is a hybrid robot, an artificial embodiment controlled by a network of living neurons and




Fig. 2. Hybrot: hybrid neural-robotics for neuroscience research. A living neuronal network is cultured on a multi-electrode array (MEA) where its activity is recorded, processed in real time, and used to control a robotic or simulated embodiment, such as the K-Team Khepera or Koala (pictured at lower right). The robot's input from proximity sensors is converted to electrical stimuli that are fed back to the neuronal network within milliseconds via a custom multi-electrode stimulation system. The hybrot's brain (MEA culture) can be imaged continuously on the microscope while its body behaves and learns. The microscope is enclosed in an incubator (lower left) to maintain the health of the living network. This closed-loop *Embodied Cultured Networks* approach may shed light on the morphological correlates of memory formation, and provide AI researchers with ideas about how to build brain-style AI.

glia cultured on a multi-electrode array (MEA) [39-41]. It will be helpful in studying some of these difficult neuroscience questions. We now have the hardware and software necessary to create a real-time loop whereby neural activity is used to control a robot, and its sensory inputs are fed back to the cultured network as patterns of electrical or chemical stimuli (Fig. 2; [42]).

These embodied cultured networks bring in vitro neuroscience models out of sensory deprivation and into the real world. They form a much needed bridge between neuroscience and AI. An MEA culture is amenable to high-resolution optical imaging [43], while the hybrot is behaving and learning, from milliseconds to months [44]. This allows correlations to be made between neural function and structure, in a living, awake and behaving subject. One of our hybrots, called MEART, was used to create portraits of viewers in a gallery. Its sensory feedback, images of its drawings in progress, affected the next action of the robotic drawing arm, in a closed-loop fashion [45]. This has been used to explore the neural mechanisms of creativity. Whether a network of a few thousand neurons can be creative is still up for debate, but it is vastly more complex than any existing artificial neural network. By studying embodied cultured networks with these new tools, we may learn some new aspects of network dynamics, memory storage, and sensory processing that could be used to make AI less artificial [41].

5 Neuroscience to AI and Back Again

Biologically-inspired artificial neural networks [46], mixed analog/digital circuits [47] and computational neuroscience approaches that attempt to elucidate brain networks [48] (as opposed to cellular properties) are gradually becoming more tightly coupled to experimental neuroscience. The fields of Psychology and Cognitive Science have traditionally made progress using theoretical foundations having little or no basis in neuroscience, due to a lingering Cartesian dualism in the thinking of their practitioners [49]. However, with neuroscience advances in psychopharmacology (e.g., more targeted neuroactive drugs) and functional brain imaging (e.g., functional MRI), Psychology and Cognitive Science advances are becoming increasingly based on and inspired by biological mechanisms. It's time for AI to move in the brainwards direction. This could involve PhD programs that merge AI and neuroscience, journals that seek to unite the two fields, and more conferences (such as the one that spawned this book) at which AI researchers and neuroscientists engage in productive dialogues.

Neuroscientists have not exactly embraced AI either. Both sides need to venture across the divide and learn what the other has to offer. How can neuroscience benefit from AI? As we have seen, brains are far too complicated for us to understand at present. AI can produce new tools for compiling the mass of neuroscience results published, and coming up with connections or general theoretical principles. On a more mundane but important level, we need AI just to help us deal with the massive data sets that modern neuroscience tools produce, such as multi-electrode arrays and real-time imaging. The new field of neuroinformatics has to date been mostly concerned with neural database design and management. Soon, with help from AI, it will incorporate more data mining, knowledge discovery, graphic visualization, segmentation and pattern recognition, and other advances yet to be invented. One can

imagine that by increasing the synergy between AI and neuroscience, a bootstrapping process will occur: more neuroscience research will inform better AI, and better AI will give neuroscientists the tools to make more discoveries and interpret them. Where it will lead, who knows, but it will be an exciting ride!

Acknowledgments. I thank Max Lungarella for insightful comments. This work was funded in part by the US National Institutes of Health (National Inst. of Mental Health, National Inst. of Neurological Disorders and Stroke, National Inst. on Drug Abuse), the US National Science Foundation Center for Behavioral Neuroscience, The Whitaker Foundation, The Keck Foundation, and the Coulter Foundation.

References

1. Lazebnik, Y.: Can a biologist fix a radio?—Or, what I learned while studying apoptosis. *Cancer Cell* 2, 179–182 (2002)
2. Purves, D., Augustine, G.J., Fitzpatrick, D., Hall, W.C., Lamantia, A.S., McNamara, J.O., et al.: *Neuroscience*, 3rd edn. Sinauer Associates, New York (2004)
3. Bear, M.F., Connors, B., Paradiso, M.: *Neuroscience: Exploring the Brain*. Lippincott Williams & Wilkins, New York (2006)
4. Kandel, E.R., Schwartz, J.H., Jessel, T.M.: *Principles of Neural Science*, 4th edn. McGraw-Hill Publishing Co., New York (2000)
5. von Neumann, J.: *The Computer and the Brain*. Yale University Press, New Haven (1958)
6. Tononi, G., Sporns, O., Edelman, G.M.: A Measure for Brain Complexity: Relating Functional Segregation and Integration in the Nervous System. *PNAS* 91(11), 5033–5037 (1994)
7. Kosslyn, S.M., Pascual-Leone, A., Felician, O., Camposano, S., Keenan, J.P., Thompson, W.L., et al.: The role of area 17 in visual imagery: convergent evidence from PET and rTMS. *Science* 284(5411), 167–170 (1999)
8. Alberini, C.M., Milekic, M.H., Tronel, S.: Mechanisms of memory stabilization and de-stabilization. *Cell Mol. Life Sci.* 63(9), 999–1008 (2006)
9. Bell, A.J.: Levels and loops: the future of artificial intelligence and neuroscience. *Philosophical Transactions of the Royal Society of London Series B-Biological Sciences* 354(1392), 2013–2020 (1999)
10. Buzsaki, G.: *Rhythms of the Brain*. Oxford U. Press, Oxford (2006)
11. Wehr, M., Laurent, G.: Odour encoding by temporal sequences of firing in oscillating neural assemblies. *Nature* 384(6605), 162–166 (1996)
12. Yuste, R., Majewska, A.: On the function of dendritic spines. *Neuroscientist* 7(5), 387–395 (2001)
13. Liu, S.-C., Kramer, J., Indiveri, G., Delbrück, T., Douglas, R.: *Analog VLSI: Circuits and Principles*. MIT Press, Cambridge (2002)
14. Reich, D.S., Mechler, F., Purpura, K.P., Victor, L.D.: Interspike intervals, receptive fields, and information encoding in primary visual cortex. *Journal of Neuroscience* 20(5), 1964–1974 (2000)
15. Froemke, R.C., Dan, Y.: Spike-timing-dependent synaptic modification induced by natural spike trains. *Nature* 416(6879), 433–438 (2002)
16. Lungarella, M., Pegors, T., Bulwinkle, D., Sporns, O.: Methods for Quantifying the Informational Structure of Sensory and Motor Data. *Neuroinformatics* 3, 243–262 (2005)

17. Lungarella, M., Sporns, O.: Mapping Information Flow in Sensorimotor Networks. *PLOS Computational Biology* 2, 1301–1312 (2006)
18. Mehta, S.B., Whitmer, D., Figueroa, R., Williams, B.A., Kleinfeld, D.: Active Spatial Perception in the Vibrissa Scanning Sensorimotor System. *PLoS Biol.* 5(2), e15 (2007)
19. Chiel, H.J., Beer, R.D.: The brain has a body: adaptive behavior emerges from interactions of nervous system, body and environment. *Trends in Neurosciences* 20(12), 553–557 (1997)
20. Lipson, H.: Curious and Creative Machines. In: *The 50th Anniversary Summit of Artificial Intelligence*, Ascona, Switzerland (2006)
21. Sporns, O., Tononi, G., Edelman, G.M.: Connectivity and complexity: the relationship between neuroanatomy and brain dynamics. *Neural Networks* 13(8-9), 909–922 (2000)
22. Stoffregen, T.A., Pittenger, J.B.: Human echolocation as a basic form of perception and action. *Ecological Psychology* 7, 181–216 (1995)
23. Bach-y-Rita, P.: *Brain Mechanisms in Sensory Substitution*. Academic Press, New York (1972)
24. Dennett, D.C.: *Consciousness Explained*. Little, Brown & Co., Boston (1991)
25. Ramon y Cajal, S.: *Histologie du Systeme Nervex de l'Homme et des Vertebres*. Maloine, Paris (1911)
26. Ramon y Cajal, S.: *Les nouvelles idées sur la structure du système nerveux chez l'homme et chez les vertébrés* (L. Azoulay, Trans.). C. Reinwald & Cie, Paris (1894)
27. Kalinichenko, S.G., Okhotin, V.E.: Unipolar brush cells—a new type of excitatory interneuron in the cerebellar cortex and cochlear nuclei of the brainstem. *Neurosci. Behav. Physiol.* 35(1), 21–36 (2005)
28. Sporns, O., Tononi, G., Kötter, R.: The Human Connectome: A Structural Description of the Human Brain. *PLOS Computational Biology* 1, 245–251 (2005)
29. Watts, D.J., Strogatz, S.H.: Collective dynamics of 'small-world' networks. *Nature* 393(6684), 440–442 (1998)
30. Gerstner, W., Kreiter, A.K., Markram, H., Herz, A.V.M.: Neural codes: Firing rates and beyond. *Proc. Natl. Acad. Sci. USA* 94, 12740–12741 (1997)
31. Bi, G.Q., Poo, M.M.: Synaptic modifications in cultured hippocampal neurons: Dependence on spike timing, synaptic strength, and postsynaptic cell type. *Journal of Neuroscience* 18, 10464–10472 (1998)
32. Izhikevich, E.M.: Polychronization: Computation with spikes. *Neural Computation* 18(2), 245–282 (2006)
33. Maass, W., Natschlager, T., Markram, H.: Real-time computing without stable states: a new framework for neural computation based on perturbations. *Neural Comput.* 14(11), 2531–2560 (2002)
34. Markram, H.: The blue brain project. *Nat. Rev. Neurosci.* 7(2), 153–160 (2006)
35. Mel, B.W.: Information-processing in dendritic trees. *Neural Computation* 6, 1031–1085 (1994)
36. Weiler, I.J., Hawrylak, N., Greenough, W.T.: Morphogenesis in Memory Formation - Synaptic and Cellular Mechanisms. *Behavioural Brain Research* 66(1-2), 1–6 (1995)
37. Majewska, A., Sur, M.: Motility of dendritic spines in visual cortex in vivo: Changes during the critical period and effects of visual deprivation. *Proc. Natl. Acad. Sci. USA* 100(26), 16024–16029 (2003)
38. Leuner, B., Falduto, J., Shors, T.J.: Associative Memory Formation Increases the Observation of Dendritic Spines in the Hippocampus. *J. Neurosci.* 23(2), 659–665 (2003)

39. DeMarse, T.B., Wagenaar, D.A., Potter, S.M.: The neurally-controlled artificial animal: A neural-computer interface between cultured neural networks and a robotic body. *Society for Neuroscience Abstracts* 28, 347.1 (2002)
40. DeMarse, T.B., Wagenaar, D.A., Blau, A.W., Potter, S.M.: The neurally controlled animat: Biological brains acting with simulated bodies. *Autonomous Robots* 11, 305–310 (2001)
41. Bakkum, D.J., Shkolnik, A.C., Ben-Ary, G., Gamblen, P., DeMarse, T.B., Potter, S.M.: Removing some ‘A’ from AI: Embodied Cultured Networks. In: Iida, F., Pfeifer, R., Steels, L., Kuniyoshi, Y. (eds.) *Embodied Artificial Intelligence*. LNCS (LNAI), vol. 3139, pp. 130–145. Springer, Heidelberg (2004)
42. Potter, S.M., Wagenaar, D.A., DeMarse, T.B.: Closing the Loop: Stimulation Feedback Systems for Embodied MEA Cultures. In: Taketani, M., Baudry, M. (eds.) *Advances in Network Electrophysiology using Multi-Electrode Arrays*, pp. 215–242. Springer, New York (2006)
43. Potter, S.M.: Two-photon microscopy for 4D imaging of living neurons. In: Yuste, R., Konnerth, A. (eds.) *Imaging in Neuroscience and Development: A Laboratory Manual*, pp. 59–70. Cold Spring Harbor Laboratory Press (2005)
44. Potter, S.M., DeMarse, T.B.: A new approach to neural cell culture for long-term studies. *J. Neurosci. Methods* 110, 17–24 (2001)
45. Bakkum, D.J., Chao, Z.C., Gamblen, P., Ben-Ary, G., Potter, S.M.: Embodying Cultured Networks with a Robotic Drawing Arm. In: *The 29th IEEE EMBS Annual International Conference* (2007)
46. Granger, R.: Engines of the brain: the computational instruction set of human cognition. *AI Magazine* 27, 15 (2006)
47. Linares-Barranco, A., Jiminez-Moreno, G., Linares-Barranco, B., Civit-Balcells, A.: On Algorithmic Rate-Coded AER Generation. *IEEE Transactions on Neural Networks* 17, 771–788 (2006)
48. Seth, A.K., McKinsty, J.L., Edelman, G.M., Krichmar, J.L.: Visual binding through reentrant connectivity and dynamic synchronization in a brain-based device. *Cerebral Cortex* 14, 1185–1199 (2004)
49. Damasio, A.R.: *Descartes’ Error: Emotion, Reason, and the Human Brain*. Gosset/Putnam Press, New York (1994)