

How Brains Are Built: Principles of Computational Neuroscience
By Richard Granger, Ph.D.

Editor's note: The goal of computational neuroscience is to understand the brain and its mechanisms well enough to artificially simulate their functions. In some areas, like hearing, vision, and prosthetics, there have been great advances in the field. Yet there is still much about the brain that is unknown and therefore cannot be artificially replicated: How does the brain use language, make complex associations, or organize learned experiences? Once the neural pathways responsible for these and many other functions are fully understood and reconstructed, we will have the ability to build systems that can match—and maybe even exceed—the brain's capabilities.

Article available online at <http://dana.org/news/cerebrum/detail.aspx?id=30356>

“If I cannot build it, I do not understand it.” So said Nobel laureate Richard Feynman, and by his metric, we understand a bit about physics, less about chemistry, and almost nothing about biology.¹

When we fully understand a phenomenon, we can specify its entire sequence of events, causes, and effects so completely that it is possible to fully simulate it, with all its internal mechanisms intact. Achieving that level of understanding is rare. It is commensurate with constructing a full design for a machine that could serve as a stand-in for the thing being studied. To understand a phenomenon sufficiently to fully simulate it is to understand it *computationally*.

“Computation” does not refer to computers per se; rather it refers to the underlying principles and methods that make them work. As Turing Award recipient Edsger Dijkstra said, computational science “is no more about computers than astronomy is about telescopes.”² Computational science is the study of the hidden rules underlying complex phenomena from physics to psychology.

Computational neuroscience, then, has the aim of understanding brains sufficiently well to be able to simulate their functions, thereby subsuming the twin goals of science and engineering: deeply understanding the inner workings of our brains, and being able to construct simulacra of them. As simple robots today substitute for human physical abilities, in settings from factories to hospitals, so brain engineering will construct stand-ins for our mental abilities—and possibly even enable us to fix our brains when they break.

Brains and Their Construction

Brains, at one level, consist of ion channels, chemical pumps, specialized proteins. At another level, they contain several types of neurons connected via synaptic junctions. These are in turn composed into networks consisting of repeating modules of carefully arranged circuits. These networks are arrayed in interacting brain structures and systems, each with distinct internal wiring and each carrying out distinct functions. As in most complex systems, each level arises from those below it but is not readily reducible to its constituents. Our understanding of an organism depends on our understanding of its component organs, but also on the ongoing interactions among those parts, as is evident in differentiating a living organism from a dead one.

For instance, kidneys serve primarily to separate and excrete toxins from blood and to regulate chemical balances and blood pressure, so a kidney simulacrum would entail a nearly complete set of chemical and enzymatic reactions. A brain also monitors many critical regulatory mechanisms, and a complete understanding of it will include detailed chemical and biophysical characteristics.

But brains, alone among organs, produce thought, learning, recognition. No amount of engineering has yet equaled, let alone surpassed, brains’ abilities at these tasks. Despite huge efforts and

large budgets, we have no artificial systems that rival humans at recognizing faces, nor understanding natural languages, nor learning from experience.

There are, then, crucial principles that brains encode that have so far eluded the best efforts of scientists and engineers to decode. Much of computational neuroscience is aimed directly at attempting to decipher these principles.

Today we cannot yet fully simulate every aspect of a kidney, but we have passed a decisive threshold: we can build systems that replicate kidney principles so closely that they can supplant their function in patients who have suffered kidney loss or damage. Artificial kidneys do not use the same substrate as real kidneys; circuits and microfluidics take the place of cells and tissue, yet they carry out operations that are equivalent, and lifesaving, to the human bodies that use them. A primary long-term goal of computational neuroscience is to derive scientific principles of brain operation that will catalyze the comparable development of prosthetic brains and brain parts.

Do We Know Enough About Brains to Build Them?

As with any complex system, in the absence of full computational understanding of the brain, we proceed by collecting constraints: experimentally observable data can rule out potential explanations. The more we can rule out, the closer we are to hypotheses that can account for the facts. Many constraining observations have usefully narrowed our understanding of how mental activity arises from brain circuitry; these can be organized into five key categories.

Brain component allometry: Remarkably tight relationships hold between a brain's overall size and the size of its constituent components. Just knowing the overall brain size of any mammal, we can with great precision predict the size of all component structures within the brain. Thus, with few exceptions, brains apparently do not and cannot choose which structures to differentially expand or reconfigure.³⁻¹¹ So, quite surprisingly, rather than a range of different circuits, or even selective resizing of brain components, human brains are instead largely built from the same components as other mammalian brains, in the same circuit layouts, with highly predictable relative sizes. Apparently a quantitative change (brain size) results in a qualitative one (uniquely human computational capabilities).^{9, 12}

Telencephalic uniformity: Circuits throughout the forebrain (telencephalon) exhibit notably similar repeated designs,^{6,13} with few exceptions,¹⁴⁻¹⁹ including some slightly different cell types, circuit structures, and genes. Yet brain areas purported to underlie unique human abilities (e.g., language) barely differ from other structures; there are no extant hypotheses of how the modest observed genetic or anatomical differences could engender exceedingly different functions. Taken together, these findings

intimate the existence of a few elemental core computational functions that are re-used for a broad range of apparently different sensory and cognitive operations.

Anatomical and physiological imprecision: Evidence suggests that neural components are surprisingly sloppy (probabilistic) in their operation, very sparsely connected, low-precision, and extraordinarily slow,²⁰⁻²² despite exhibiting careful timing under some experimental conditions.²³⁻²⁷ Either brains are far more precise than we yet understand, or else they carry out families of algorithms whereby precise computations arise from imprecise components.²⁸⁻³¹ If so, this greatly constrains the types of operations that any brain circuits could be engaged in.

Task specification: Though artificial telephone operators field phone inquiries with impressive voice recognition, we know that they could do far better. The only reason we know this is that human operators substantially outperform them; there are no other formal specifications whatsoever that characterize the voice recognition task.^{32,33} Engineers began by believing that they understood the task sufficiently to construct artificial operators. It has turned out that their specification of the task does not match the actual, still highly elusive set of steps that humans actually perform in recognizing speech. Without formal task specifications, the only way to equal human performance may be to come to understand the brain mechanisms that give rise to the behavior.

Parallel processing: Some recognition tasks take barely a few hundred milliseconds,^{34,35} corresponding to no more than hundreds of serial neural steps (of milliseconds each), strongly indicating myriad neurons acting in parallel,³⁶ imposing a very strong constraint on the types of operations that individual neurons could be carrying out. Yet parallelism in computer science, even on a small scale, such as two or three simultaneous operations, has proven very elusive. Why, for instance, don't our dual-core or quad-core computers run two or four times faster than single-core systems? The (painfully direct) answer is that we simply do not yet know how to divide most software into parts that can effectively exploit the presence of these additional hardware elements. Even for readily parallelizable software, it is challenging to design hardware that yields scalable returns as processors are added.^{37,38} It is increasingly possible that principles of brain architecture may help identify novel and powerful parallel machine designs.

From Circuits to Algorithms to Prosthetics

There are several promising instances in which different laboratories (even laboratories that are competing with each other) have arrived at substantial points of agreement about what certain brain areas are likely doing. A notable success story arises from studies of the basal ganglia, which takes two kinds of

inputs: sensory information from the neocortex, and “reward” and “punishment” information arising from external stimuli. We are close to computationally understanding this large chunk of the brain, which apparently carries out just one of our primary learning abilities: our slow “trial and error” learning (studied in computational neuroscience as “reinforcement learning”), underlying our ability to acquire such skills as riding a bike.^{30, 39–65}

In addition, there is a growing consensus that circuits in the neocortex, by far the largest set of brain structures in humans, carry out another, quite different kind of learning: the ability to rapidly learn new facts and to organize newly acquired knowledge into vast hierarchical structures that encode complex relationships, such as categories and subcategories, episodes, and relations.^{28, 66–74}

And these two systems are connected to each other, via far-reaching cortico-basal ganglia (aka cortico-striatal) loops.⁴⁹ The basal ganglia system carries out the computational operations of skill learning (reinforcement learning) while cortical circuits computationally construct vast hierarchies of facts and relations among facts. Interestingly, computational research on reinforcement learning has found that adding hierarchies to the process can greatly improve learning performance.^{75,76} Our ancestors (reptiles and early mammals) were largely driven by the basal ganglia, whereas mammalian evolution has hugely expanded the relative size of the neocortex. By consistently increasing the size ratio of the neocortex to the basal ganglia, mammalian brain evolution may be solving a specific computational puzzle.^{29, 75–79} Our understanding of human and animal learning abilities is being advanced by these computational studies, and we are developing novel methods for machine learning, enabling more powerful computer algorithms for analysis of complex data ranging from medical to commercial to financial applications.

Meanwhile, as study of these primary cortico-striatal brain structures remains very much still in progress, great advances have been made in deep, computational understanding of certain circumscribed brain systems, in particular those involved in early sensory transduction and perception. The results have been striking.

Analysis of cochlear mechanisms has led to the construction of prosthetics that serve today as cures for more than 100,000 people who have lost their hearing.⁸⁰ Retinal prosthetics are in advanced development.^{81–85} In a recent study, patients with retinal implants recognized printed letters of size and distance comparable to reading a book in relatively low light. And experimental prosthetic arms can respond to brain-initiated control; people learn to control the arm simply by deciding to move it.^{86, 87}

These sensory and motor findings have also led to formalizations of the general problem of acting in environments that are only partly observable and are dynamically changing, such as robotics or automated navigation; the result is a set of increasingly impressive robotic methods that see and navigate in complex surroundings.⁸⁸ In a series of trials run by the Department of Defense over the last several

years, vehicles were, for the first time, able to navigate through real urban traffic, merging, passing, parking, and negotiating intersections, with no human control. Retinal algorithms operate equally well on other sensors such as radar; and prosthetic limb algorithms are wholly applicable to robots. Many of the algorithms that operate robots and automated vehicles are closely related to those that operate prosthetic limbs.

As we come to computationally understand how these peripheral sensorimotor systems work, the distinction between natural and artificial is being eroded. A breed of robots that share many of our own dexterity and perceptual abilities is likely to emerge directly from this research. As these increasingly biologically-based robots, or biots, come to replace human skilled labor, the economic and social consequences may be substantial.

From Percept to Concept

The primary differences between human brains and those of other animals lie not in our sensory or motor mechanisms, which are largely shared across many species, but rather in cognitive abilities: association, representation, reasoning. Despite great advances in peripheral prosthetics, there is no commensurate understanding of advanced cognition.

The abilities of peripheral circuits (retina, cochlea, initial thalamic and cortical regions) are largely built in at birth via genetic programs and shaped in early childhood during developmentally critical periods. In contrast, the rest of the neocortex will use those built-in systems to acquire masses of specific information about the environment over a lifetime. Neocortical circuits are not born with knowledge of particular scenes, faces, or actions; these are acquired through sensorimotor experience: observing and interacting with objects and events in our surroundings. Cortical circuits are engaged almost entirely in fact learning: rapid, permanent acquisition and organization of everyday occurrences. The low-level biological mechanisms underpinning long-term fact learning (permanent, anatomical synaptic changes, rather than inherently ephemeral chemical changes) are becoming understood.⁸⁹ But the neocortex is not just a passive warehouse of billions of isolated facts; we can arbitrarily associate them, recall them, embellish them.³³ Association, recall, retrieval, organization—all that we can actually do with memory—depends on mechanisms that are as yet still unknown.

Early cortical areas, then, deal with recognizing objects (even in different lighting, settings, and clutter), but some laboratories are increasingly focusing on cortical circuits that are beyond the early sensory areas: the vast remainder of the neocortex that somehow encodes sequences, associations, and abstract relations.^{33, 90–99}

Seeing a phone, we perceive not only its visual form but also its affordances (calling, texting, photographing, playing music), our memories of it (when we got it, where we have recently used it), and a

wealth of potential associations (our ringtone, whom we might call, whether it is charged, etc.). The questions of how cross-modal information is learned and integrated, and in what form the knowledge is stored—how percepts become concepts—now constitute the primary frontier of work in computational neuroscience. In this borderland between perception and cognition, the peripheral language of the senses is transmuted to the internal lingua franca of the brain, freed from literal sensation and formulated into internal representations that can include a wealth of associations.

Even our simplest perceptions often rely on top-down processing: using stored memory representations to inform our ongoing perception and recognition. In some circumstances, we can recognize objects in just tens of milliseconds,^{34,35} so rapidly that it is unlikely that any top-down pathways are yet engaged. Yet once we're beyond simple recognition, to the far richer range of inference, association, and even language, memories strongly influence our perceptions. Merely thinking of a car is sufficient to activate the same early visual areas that would have been triggered by actually seeing the car, including its shape, size, color, and other features.^{100–102}

These early visual areas are just one instance of the spread of activation from a triggering memory.^{103–105} Thinking of a car may also activate many other areas, as yet largely unmapped, that encode knowledge of how to open car doors, turn ignition keys, steer, accelerate, brake—or information about what particular car you own, where it is parked, and so on. Today we can experimentally test for visual shape information because we know a great deal about how to decode neural responses that occur in early visual areas,¹⁰⁶ but we have comparatively modest data for other associative knowledge.^{107–109} Computational models of spreading activation^{110,111} are now striving to make contact with specific neural mechanisms and brain pathways, to arrive at convergent hypotheses like those of peripheral sensory systems.

Computing Individual Differences: From Neurotypes to Cognotypes

Though all of us have extraordinarily similar brains, even small differences can be striking. Whether particular characteristics are genetic, developmental, or learned is still often impossible to ascertain, but individual behavioral differences are highly likely to directly correspond to individual brain differences, whether genetic or acquired. Most work in computational neuroscience—from perception to cognition, from anatomy to computational models—has focused on one agent at a time, one brain at a time. A further frontier will be to confront differences among individuals.

Our bodies are built by genetic programs that became locked into particular patterns early on in mammalian evolution: four appendages; eyes above nose above mouth between ears; ten fingers and ten toes. We are not optimized to have just these features and no others; most of the variations that we might imagine—nose above eyes; five limbs; tentacles instead of hands—have never been tried by evolution,

not because they're not improvements, but simply because they may not be consistent with the genetic prebuilt modules that have been bundled for hundreds of millions of years.^{112,113}

Brain components are body components, so it is not surprising that evolutionary brain changes also are highly predictable, exhibiting selectional pressure only within the constraints of prescribed regularities: all mammals have almost exactly the same brain regions, in the same allometric size relationship, wired extraordinarily similarly.^{5,6,9} It is hypothesized that the relatively modest brain differences that do occur are “canalized” into a relatively small set of categories. Brains create behaviors, and brain differences can create behavioral differences.

Because brain differences tend to follow certain patterns of architectural arrangements, or neurotypes, individual differences then tend to fall into groups, which correspondingly can be referred to as cognotypes. These can be described as a range of recognized characteristics of differential cognitive types,^{114,115} such as types of psychopathy, personality attributes of introversion or extraversion, Asperger's characteristics, even different motor skill abilities. Any of these may entail a complex combination of inherited (genetic or polygenic) and acquired (developed or learned) characteristics. Symptoms of Asperger's, for instance, can include the seemingly arbitrary combination of high mathematical and engineering abilities with low social abilities, whereas there tend not to be behavioral types combining, say, high empathy with synesthesia, or low motor abilities with unusually high face-recognition abilities.

We won't achieve a full computational understanding of human brains until we understand why only certain variants tend to occur, and until we can model how it is that architectural brain differences can mechanistically generate cognitive differences. There are likely to be salient philosophical questions of will and intent, and ethical questions of capacity and culpability that will, it is hoped, be clarified as our understanding deepens.

Extrapolations

The field of computational neuroscience is uniquely situated to decrypt the brain's mechanisms, to construe the perceptual and memorial abilities that still stymie our best engineering efforts. Once we crack the code, we finally will be able to construct systems that equal human performance at perceptual tasks. And having finally understood the underlying mechanisms, we may very well be able at long last to improve on them. There is no known formal reason why the capabilities of our brains may not eventually be equaled or exceeded.

There are economics to these advances, and policy implications abound. When auditory implants first became available, the scientific community widely doubted their efficacy. It took years of demonstration before they were accepted. They are expensive: today their cost can run to \$100,000 per

patient. And there are risks: the surgical implantation procedure may lead to a higher incidence of meningitis.^{116,117} Moreover, there are social complications: some in the deaf community find cochlear implants to be ethically misplaced, arguing that the deaf should not be thought of as disabled at all, but rather as a “minority cultural group.”¹¹⁸

What of brain parts that are deeper than just the peripheral hearing system? Traumatic brain injury can cause debilitating deficits in memory and cognition; at present, such injuries are extremely difficult even to diagnose, let alone to treat. Implants to restore lost cognitive abilities for such accident victims would be revolutionary, and would be welcomed.

But if implants existed for accident-induced cognitive losses, could they also be used to augment uninjured cognitive function? There is suggestive evidence from drugs: some Alzheimer’s medications may improve memory in people with mild cognitive impairment—but the FDA has not yet approved the use of any treatments for these lesser conditions.^{119,120} How would regulators at the FDA react if it became possible to augment our brains—implants to help us think faster or to increase our memory capacity? The economic, social, and political concomitants of such technology would surely eclipse those arising from cochlear implants.

Each brain contains idiosyncrasies; our brains define who we are. The way we interact, the kinds of decisions we make, the connections we perceive—all arise from the still-obscure mechanisms of the vast span of thalamocortical circuits and cortico-striatal loops in our heads. These repeating components give us our mammalian abilities, our uniquely human faculties, and our individual characteristics. The computational understanding of individual and group differences will likely lead to a new science of different types of cognitive behavior, with implications ranging from law to education. The formerly familiar terrain of human nature may appear quite different in this light; perhaps, arriving there, we will truly know the place for the first time.

Our abilities are not inimitable; brain circuits are circuits, albeit nonstandard ones, and they will yield to analysis. As computational neuroscience comes to demystify them, we verge on an era of new frontiers in science and medicine, in which we can increasingly repair, enhance, and likely supplant the biological engines we think with.

Richard Granger, Ph.D., is a professor at Dartmouth with faculty positions in the departments of psychological and brain sciences, computer science, and the Thayer School of Engineering. He directs Dartmouth's interdisciplinary Brain Engineering Laboratory, with research projects ranging from computation and robotics to neuroimaging and cognitive neuroscience. He has authored more than 100 scientific papers and holds numerous issued patents, is an elected fellow of the American Association for the Advancement of Science (AAAS), and serves on the boards of a number of technology corporations and government agencies. He is co-inventor of FDA-approved devices and drugs in clinical trials, and has been the principal architect of a series of advanced computational systems for military, commercial, and medical applications.

References

1. Feynman, R. In Hawking, S. (2001). *The universe in a nutshell* (p. 83). Bantam.
2. Dijkstra, E. (2001). *Denken als Discipline (Discipline in Thought)* [Video interview]. Retrieved from <http://www.cs.utexas.edu/users/EWD/video-audio/NoorderlichtVideo.html>.
3. Jerison, H. (1973). *Evolution of the brain and intelligence*. Academic Press.
4. Finlay, B., Innocenti, G., & Scheich, H. (1991). *The neocortex: Ontogeny and phylogeny*. Plenum Press.
5. Finlay, B., & Darlington, R. (1995). Linked regularities in the development and evolution of mammalian brains. *Science*, 268, 1578–1584.
6. Striedter, G. F. (2005). *Principles of brain evolution*. Sinauer Associates.
7. Falk, D., & Gibson, K. (2001). *Evolutionary anatomy of the primate cerebral cortex*. Cambridge University Press.
8. Sherwood, C., Holloway, R., Semendeferi, K., & Hof, P. (2010). Inhibitory interneurons of the human prefrontal cortex display conserved evolution of the phenotype and related genes. *Proceedings of the Royal Academy of Science B*, 277, 1011–1020.
9. Lynch, G., & Granger, R. (2008). *Big brain*. Palgrave Macmillan.
10. Herculano-Houzel, S. (2009). The human brain in numbers: A linearly scaled-up primate brain. *Frontiers in Neuroscience*, 3, 1–11.
11. Semendeferi, K., Teffer, K., Buxhoeveden, D., Park, M., Bludau, S., Amunts, K., . . . Buckwalter, J. (2010). Spatial organization of neurons in the prefrontal cortex sets humans apart from great apes. *Cerebral Cortex*. doi: 10.1093/cercor/bhq191.
12. Amati, D., & Shallice, T. (2007). On the emergence of modern humans. *Cognition*, 103(3), 358–385.

13. Jones, E. G., & Rakic, P. (2010). Radial columns in cortical architecture: It is the composition that counts. *Cerebral Cortex*, 20(10), 2261–2264.
14. Nimchinsky, E., Glissen, E., Allman, J., Perl, D., Erwin, J., & Hof, P. (1999). A neuronal morphologic type unique to humans and great apes. *Proceedings of the National Academy of Science*, 96, 5268–5273.
15. Galuske, R., Schlote, W., Bratzke, H., & Singer, W. (2000). Interhemispheric asymmetries of the modular structure in humans. *Science*, 289, 1946–1949.
16. Buxhoeveden, D., Switala, A., Roy, E., Litaker, M., & Casanova, M. (2001). Morphological differences between minicolumns in human and nonhuman primate cortex. *American Journal of Physical Anthropology*, 115, 361–371.
17. Lai, C., Fisher, S., Hurst, J., Levy, E., Hodgson, S., Fox, M., . . . Monaco, A. (2000). The SPCH1 region on human 7q31: Genomic characterization of the critical interval and localization of translocations associated with speech and language disorder. *American Journal of Human Genetics*, 67, 357–368.
18. Evans, P., Gilbert, S., Mekel-Bobrov, N., Vallender, E., Anderson, J., Vaez-Azizi, L., . . . Lahn, B. (2005). Microcephalin, a gene regulating brain size, continues to evolve adaptively in humans. *Science*, 309, 1717–1720.
19. Mekel-Bobrov, N., Gilbert, S., Evans, P., Vallender, E., Anderson, J., Hudson, R., . . . Lahn, B. (2005). Ongoing adaptive evolution of ASPM, a brain size determinant in *Homo sapiens*. *Science*, 309, 1720.
20. Braitenberg, V., & Schüz, A. (1998). *Cortex: Statistics and geometry of neuronal connectivity*. Springer-Verlag.
21. Häusser, M., & Mel, B. (2003). Dendrites: Bug or feature? *Current Opinion in Neurobiology*, 13, 372–383.
22. Fuhrmann, G., Segev, I., Markram, H., & Tsodyks, M. (2002). Coding of temporal information by activity-dependent synapses. *Journal of Neurophysiology*, 87, 140–148.
23. Singer, W. (1999). Neuronal synchrony: A versatile code for the definition of relations? *Neuron*, 24(1), 49–65, 111–125.
24. Singer, W. (2010). Distributed processing and temporal codes in neuronal networks. *Cognitive Neurodynamics* 3, 189–196.
25. Traub, R., Bibbig, A., LeBeau, F., Buhl, E., & Whittington, M. (2004). Cellular mechanisms of neuronal population oscillations in the hippocampus in vitro. *Annual Review of Neuroscience*, 27, 247–278.
26. Wang, X. (2010). Neurophysiological and computational principles of cortical rhythms in cognition. *Physiological Reviews*, 90(3), 1195–1268.
27. Clopath, C., Busing, L., Vasilaki, E., & Gerstner, W. (2010). Connectivity reflects coding: A model of voltage-based STDP with homeostasis. *Nature Neuroscience*, 13, 344–352.

28. Rodriguez, A., Whitson, J., & Granger, R. (2004). Derivation and analysis of basic computational operations of thalamocortical circuits. *Journal of Cognitive Neuroscience*, 16, 856–877.
29. Granger, R. (2005). Brain circuit implementation: High-precision computation from low-precision components. In Berger & Glanzman (Eds.), *Replacement parts for the brain* (pp. 277–294). MIT Press.
30. Granger, R. (2006). Engines of the brain: The computational instruction set of human cognition. *AI Magazine*, 27, 15–32.
31. Felch, A., & Granger, R. (2008). The hypergeometric connectivity hypothesis: Divergent performance of brain circuits with different synaptic connectivity distributions. *Brain Research*, 1202, 3–13.
32. Edelman, S. (1999). *Representation and recognition in vision*. MIT Press.
33. Edelman, S., & Intrator, N. (2003). Towards structural systematicity in distributed, statically bound visual representations. *Cognitive Science*, 27, 73–110.
34. Thorpe, S., Fize, D., & Marlot, C. (1996). Speed of processing in the human visual system. *Nature*, 381(6582), 520–522.
35. Stanford, T., Shankar, S., Massoglia, D., Costello, M., & Salinas, E. (2010). Perceptual decision making in less than 30 milliseconds. *Nature Neuroscience*, 13(3), 379–385.
36. Feldman, J., & Ballard, D. (1982). Connectionist models and their properties. *Cognitive Science*, 6, 205–254.
37. Asanovic, K., Bodik, R., Demmel, J., Keaveny, T., Keutzer, K., Kubitowicz, J., . . . Yelick, K. (2009). A view of the parallel computing landscape. *Communications of the ACM*, 52, 56–67.
38. Moorkanikara, J., Felch, A., Chandrashekar, A., Dutt, N., Granger, R., Nicolau, A., & Veidenbaum, A. (2009). Brain-derived vision algorithm on high-performance architectures. *International Journal of Parallel Programming*, 37, 345–369.
39. Schultz, W., Dayan, P., & Montague, R. (1997). A neural substrate of prediction and reward. *Science*, 175, 1593–1599.
40. Schultz, W., Apicella, P., & Ljungberg, T. (1993). Responses of monkey dopamine neurons to reward and conditioned stimuli during a delayed response task. *Journal of Neuroscience*, 13, 900–913.
41. Schultz, W. (1998). Predictive reward signal of dopamine neurons. *Journal of Neurophysiology*, 80, 1–27.
42. Schultz, W. (2002). Getting formal with dopamine and reward. *Neuron*, 36, 241–263.
43. Suri, R., & Schultz, W. (2001). Temporal difference model reproduces anticipatory neural activity. *Neural Computation*, 13(4), 841–862.

44. Suri, R. (2001). Anticipatory responses of dopamine neurons and cortical neurons reproduced by internal model. *Experimental Brain Research*, 140(2), 234–240.
45. Sutton, R. S., & Barto, A. G. (1990). Time-derivative models of Pavlovian reinforcement. In Gabriel & Moore (Eds.), *Learning and computational neuroscience: Foundations of adaptive networks*, (pp. 497–537). MIT Press.
46. Sutton, R., & Barto, A. (1998). *Reinforcement learning: An introduction*. MIT Press.
47. Strick, P., Dum, R., & Mushiake, H. (1995). Basal ganglia loops with the cerebral cortex. In M. Kimura & A. Graybiel (Eds.), *Functions of the cortico-basal ganglia loop* (pp. 106–124). Springer-Verlag.
48. Gerfen, C., & Wilson, C. (1996). The basal ganglia. In Swanson, Bjorklund, & Hokfelt (Eds.), *Handbook of Chemical Neuroanatomy*, vol. 12 (pp. 371–468). Elsevier.
49. Alexander, G., & DeLong, M. (1985). Microstimulation of the primate neostriatum. I. Physiological properties of striatal microexcitable zones. *Journal of Neurophysiology*, 53, 14001–11416.
50. Graybiel, A., Aosaki, T., Flaherty, A., & Kimura, M. (1994). The basal ganglia and adaptive motor control. *Science*, 265, 1826–1831.
51. Graybiel, A. (1995). Building action repertoires. *Current Opinion in Neurobiology*, 5, 733–741.
52. Houk, J., Davis, J., & Beiser, D. (1995). *Models of information processing in the basal ganglia*. MIT Press.
53. Houk, J., & Wise, S. (1995). Distributed modular architectures linking basal ganglia, cerebellum, and cerebral cortex. *Cerebral Cortex*, 2, 95–110.
54. Knowlton, B., & Squire, L. (1993). The learning of categories: Parallel brain systems for item memory and category knowledge. *Science*, 262, 1747–1749.
55. Brucher, F. (2000). *Reward-based learning and basal ganglia: A biologically realistic, computationally explicit theory*. Unpublished doctoral dissertation, University of California.
56. Poldrack, R., Clark, J., Pare-Blagoev, E., Shohamy, D., Cresco Moyano, J., Myers, C., & Gluck, M. (2001). Interactive memory systems in the human brain. *Nature*, 414, 546–550.
57. Daw, N. (2003). *Reinforcement learning models of the dopamine system and their behavioral implications*. Unpublished doctoral dissertation, Carnegie Mellon University.
58. Frank, M. (2005). *Dynamic dopamine modulation in the basal ganglia: A neurocomputational account of cognitive deficits in medicated and non-medicated Parkinsonism*. *Journal of Cognitive Neuroscience*, 17, 51–72.
59. Laubach, M. (2005). Who's on first? What's on second? The time course of learning in corticostriatal systems. *Trends in Neuroscience*, 28, 509–511.

60. Yin, H., & Knowlton, B. (2006). The role of the basal ganglia in habit formation. *Nature Reviews Neuroscience*, 7(6), 464–476.
61. Swinehart, C., & Abbott, L. (2006). Dimensional reduction for reward-based learning. *Network: Computation in Neural Systems*, 17(3), 235–252.
62. Hazy, T., Frank, M., & O'Reilly, R. (2007). Towards an executive without a homunculus: Computational models of the prefrontal cortex/basal ganglia system. *Philosophical Transactions of the Royal Society B*, 362, 1601–1613.
63. Green, C., Pouget, A., & Bavelier, D. (2010). Improved probabilistic inference as a general learning mechanism with action video games. *Current Biology*, 20, 1573–1579.
64. Erickson, K., Boot, W., Basak, C., Neider, M., Prakash, R., Voss, M., Graybiel, A., . . . Kramer, A. (2010). Striatal volume predicts level of video game skill acquisition. *Cerebral Cortex* doi:10.1093/cercor/bhp293.
65. Samson, R., Frank, M., & Fellous, J. (2010). Computational models of reinforcement learning: The role of dopamine as a reward signal. *Cognitive Neurodynamics*, 4, 91–105.
66. Olshausen, B. (1996). Emergence of simple cell receptive field properties by learning a sparse code for natural images. *Nature*, 381, 607.
67. Douglas, R., & Martin, K. (2004). Neuronal circuits of the neocortex. *Annual Review Neuroscience*, 27, 419–451.
68. Friston, K. (2008). Hierarchical models in the brain. *PLoS Computational Biology*, 4, e1000211.
69. George, D., & Hawkins, J. (2009). Towards a mathematical theory of cortical microcircuits. *PLoS Computational Biology*, 5, e1000532.
70. Riesenhuber, M. (1999). Hierarchical models of object recognition in cortex. *Nature Neuroscience*, 2, 1019.
71. Lee, T., & Mumford, D. (2003). Hierarchical bayesian inference in the visual cortex. *Journal of the Optical Society of America*, 20, 1434–1448.
72. Granger, R., & Hearn, R. (2008). Models of the thalamocortical system. *Scholarpedia*, 2(11), 1796.
73. Nikolić, D., Häusler, S., Singer, W., & Maass, W. (2009). Distributed fading memory for stimulus properties in the primary visual cortex. *PLoS Biology*, 7(12), e1000260.
74. Smale, S., Rosasco, L., Bouvrie, J., Caponnetto, A., & Poggio, T. (2009). Mathematics of the neural response. *Foundations of Computational Mathematics*, 10(1), 67–91.
75. Dietterich, T. (2000). Hierarchical reinforcement learning with the MAXQ value function decomposition. *Journal of Artificial Intelligence Research*, 13, 227–303.
76. Barto, A., & Mahadevan, S. (2003). Recent advances in hierarchical reinforcement learning. *Discrete Event Dynamic Systems*, 13(341–379).

77. Sutton, R., Precup, D., & Singh, S. (1999). Between MDPs and semi-MDPs: A framework for temporal abstraction in reinforcement learning. *Artificial Intelligence*, 112, 181–211.
78. Granger, R. (2006). The evolution of computation in brain circuitry. *Behavioral and Brain Science* 29, 17.
79. Barry, J., Kaelbling, L., & Lozano-Perez, T. (2010). Hierarchical solution of large Markov decision processes. Technical report, MIT.
80. NIDCD. (2009). Cochlear implants. Retrieved from www.nidcd.nih.gov/health/hearing/coch.asp.
81. Weiland, J., Liu, W., & Humayun, M. (2005). Retinal prosthesis. *Annual Review of Biomedical Engineering*, 7, 361–401.
82. U.S. Dept. of Energy Office of Science. (2009). Artificial retina project. Retrieved from <http://artificialretina.energy.gov>.
83. Chen, K., Yang, Z., Hoang, L., Weiland, J., Humayu, M., & Liu, W. (2010). An integrated 256-channel epiretinal prosthesis. *IEEE Journal of Solid State Circuits*, 45, 1946–1956.
84. Zhou, C., Tao, C., Chai, X., Sun, Y., & Ren, Q. (2010). Implantable imaging system for visual prosthesis. *Artificial Organs*, 34(6), 518–522.
85. Zrenner, E., Wilke, R., Bartz-Schmidt, K. U., Gekeler, F., Besch, D., Benav, H., Bruckmann, A., . . . Stett, A. (2009). Subretinal microelectrode arrays allow blind retinitis pigmentosa patients to recognize letters and combine them to words. 2nd International Conference on Biomedical Engineering and Informatics (pp. 1–4).
86. Moritz, C., Perlmutter, S., & Fetz, E. (2008). Direct control of paralysed muscles by cortical neurons. *Nature*, 456, 639–642.
87. Velliste, M., Perel, S., Spalding, M. C., Whitford, A. S., & Schwartz, A. B. (2008). Cortical control of a prosthetic arm for self-feeding. *Nature*, 453(7198), 1098–1101.
88. Thrun, S. (2000). Probabilistic algorithms in robots. *AI Magazine*, 21, 93–109.
89. Fedulov, V., Rex, C., Simmons, D., Palmer, L., Gall, C., & Lynch, G. (2007). Evidence that long-term potentiation occurs within individual hippocampal synapses during learning. *Journal of Neuroscience*, 27, 8031–8039.
90. Op de Beeck, H., Baker, C., DiCarlo, J., & Kanwisher, N. (2006). Discrimination training alters object representations in human extrastriate cortex. *Journal of Neuroscience*, 26, 13025–13036.
91. Li, N., & DiCarlo, J. (2008). Unsupervised natural experience rapidly alters invariant object representation in visual cortex. *Science*, 321, 1502–1507.
92. Wallisch, P., & Movshon, J. A. (2008). Structure and function come unglued in the visual cortex. *Neuron*, 60(2), 195–197.
93. Pinto, N., Cox, D., & DiCarlo, J. (2008). Why is real-world visual object recognition hard? *PLoS Computational Biology*, 4(1), e27.

94. Simoncelli, E. P., & Olshausen, B. A. (2001). Natural image statistics and neural representation. *Annual Review of Neuroscience*, 24, 1193–1216.
95. Cox, D., Meier, P., Oertelt, N., & DiCarlo, J. (2005). Breaking position invariant object recognition. *Nature Neuroscience*, 8, 1145–1147.
96. Geman, S. (2006). Invariance and selectivity in the ventral visual pathway. *Journal of Physiology*, 100, 212–224.
97. Yuille, A., & Kersten, D. (2006). Vision as Bayesian inference: Analysis by synthesis? *Trends in Cognitive Sciences*, 10(7), 301–308.
98. DiCarlo, J., & Cox, D. (2007). Untangling invariant object recognition. *Trends in Cognitive Sciences*, 11, 333–341.
99. Roy, J., Riesenhuber, M., Poggio, T., & Miller, E. (2010). Prefrontal cortex activity during flexible categorization. *Journal of Neuroscience*, 30, 8519–8528.
100. Kosslyn, S., Alpert, N., Thompson, W., Maljkovic, V., Weise, S., Chabris, C., . . . Buonanno, F. (1993). Visual mental imagery activates topographically organized visual cortex: PET Investigations. *Journal of Cognitive Neuroscience*, 5, 263–287.
101. Kosslyn, S., Thompson, W., Kim, I., & Alpert, N. (1995). Topographical representations of mental images in primary visual cortex. *Nature*, 378, 496–498.
102. Slotnick, S., Thompson, W., & Kosslyn, S. (2005). Visual mental imagery induces retinotopically organized activation of early visual areas. *Cerebral Cortex*, 15(10), 1570–1583.
103. Posner, M., & Snyder, C. (1975). Attention and cognitive control. In Solso (Ed.), *Information processing and cognition: The Loyola symposium* (pp. 55–85). Erlbaum.
104. Neely, J. (1977). Semantic priming and retrieval from lexical memory: Roles of inhibitionless spreading activation and limited-capacity attention. *Journal of Experimental Psychology*, 106, 226–254.
105. Ratcliff, R., & McKoon, G. (1978). Priming in item recognition: Evidence for the propositional structure of sentences. *Journal of Verbal Learning and Verbal Behavior*, 17, 403–417.
106. Kay, K., Naselaris, T., Prenger, R., & Gallant, J. (2008). Identifying natural images from human brain activity. *Nature*, 452, 352–355.
107. Naselaris, T., Prenger, R., Kay, K., Oliver, M., & Gallant, J. (2009). Bayesian reconstruction of natural images from human brain activity. *Neuron*, 63, 902–915.
108. Lee, Y., Granger, R., & Raizada, R. (2010). How categorical are brain areas processing speech? (Under review).
109. Kriegeskorte, N. (2009). Relating population code representations between man, monkey, and computational models. *Frontiers in Neuroscience*, 3, 363–373.

110. Torralba, A., & Sinha, P. (2001). Statistical context priming for object detection. *Proceedings of the International Conference on Computer Vision*, 763–770.
111. Torralba, A. (2003). Contextual priming for object detection. *International Journal of Computer Vision*, 53, 169–191.
112. Carroll, S. (2005). *Endless forms most beautiful*. W. W. Norton.
113. Garcia-Fernandez, J. (2005). The genesis and evolution of homeobox gene clusters. *Nature Reviews Genetics*, 6, 881–892.
114. Zuckerman, M., Kuhlman, D., Joireman, J., Teta, P., & Kraft, M. (1993). A comparison of three structural models for personality: The big three, the big five, and the alternative five. *Journal of Personality and Social Psychology*, 65(4), 757–768.
115. Lynam, D., & Widiger, T. (2001). Using the five-factor model to represent the DSM-IV personality disorders: An expert consensus approach. *Journal of Abnormal Psychology*, 110, 401–412.
116. U.S. Food and Drug Administration. (2002). FDA public health notification: Risk of bacterial meningitis in children with cochlear implants. Retrieved from <http://www.fda.gov/MedicalDevices/Safety/AlertsandNotices/PublicHealthNotifications/ucm064526.htm>.
117. U.S. Food and Drug Administration. (2006). FDA public health notification: Continued risk of bacterial meningitis in children with cochlear implants with a positioner beyond twenty-four months post-implantation. Retrieved from <http://www.fda.gov/MedicalDevices/Safety/AlertsandNotices/PublicHealthNotifications/ucm062104.htm>.
118. Sparrow, R. (2005). Defending deaf culture: The case of cochlear implants. *Journal of Political Philosophy*, 13(2), 135–152.
119. Petersen, R., Smith, G., Waring, S., Ivnik, R., Tangalos, E., & Kokmen, E. (1999). Mild cognitive impairment: Clinical characterization and outcome. *Annals of Neurology*, 56, 303–308.
120. Rivas-Vasquez, R., Mendez, C., Rey, G., & Carrazana, E. (2004). Mild cognitive impairment: New neurophysiological and pharmacological target. *Archives of Clinical Neuropsychology*, 19, 11–27.