THE COMPUTATIONAL BRAIN

Patricia S. Churchland
Terrence J. Sejnowski

The Computational Brain

Patricia S. Churchland and Terrence J. Sejnowski

A Bradford Book
The MIT Press
Cambridge, Massachusetts
London, England

This book was set in Palatino by Asco Trade Typesetting Ltd., Hong Kong and was printed and bound in the United States of America.

Library of Congress Cataloging-in-Publication Data

Churchland, Patricia Smith.
The computational brain / Patricia S. Churchland and Terrence J. Sejnowski.
p. cm.
"A Bradford book."
Includes bibliographical references and index.
ISBN 0-262-03188-4
1. Brain—Computer simulation. 2. Neural networks (Computer science) 3. Neural circuitry.
I. Sejnowski, Terrence J. (Terrence Joseph) II. Title.
[DNLM: 1. Brain—physiology. 2. Computer Simulation. 3. Models, Neurological.
4. Neurosciences—methods. WL 26.5 C563c]
QP356.C48 1992
612.8'2'0113—dc20
DNLM/DLC
for Library of Congress 91-28056
CIP

Series Foreword

Computational neuroscience is an approach to understanding the information content of neural signals by modeling the nervous system at many different structural scales, including the biophysical, the circuit, and the systems levels. Computer simulations of neurons and neural networks are complementary to traditional techniques in neuroscience. This book series welcomes contributions that link theoretical studies with experimental approaches to understanding information processing in the nervous system. Areas and topics of particular interest include biophysical mechanisms for computation in neurons, computer simulations of neural circuits, models of learning, representation of sensory information in neural networks, systems models of sensory-motor integration, and computational analysis of problems in biological sensing, motor control, and perception.

Terrence J. Sejnowski
Tomaso Poggio

Preface

To understand how neurons give rise to a mental life, we must know what they do, both individually as single cells and collectively as coherent systems of cells. The idea that brains are computational in nature has spawned a range of explanatory hypotheses in theoretical neurobiology. This book represents one slant on current computational research relevant to neurobiology, a slant both on the conceptual foundations and the benchmark studies. We mustered our plans for this book against the backdrop of several earlier projects: *Parallel Distributed Processing*, edited by Dave Rumelhart and Jay McClelland (1986), and *Neurophilosophy* (P. S. Churchland, 1986). Since those books went to press, much has changed. Powerful new tools for modeling neurons and circuits of neurons have become available, and a conceptual framework for neurocomputational projects has been steadily greening out. Puzzling questions abound on every side, however, concerning such matters as algorithms for weight-setting in neural nets and the extent to which they can be valuable in neuromodeling; concerning biological realism in neural net models and what degree of realism is necessary to make a model useful; and highly focused questions such as what exactly is "Hebbian learning" and what are "grandmother" cells.

The questions that became pivotal in *The Computational Brain* were questions that have been biting our heels more or less incessantly. The book is thus shaped by what has bothered or beguiled us, individually and jointly. We learned a great deal from the conversations in the laboratory, some of which extended over many months. Francis Crick launched the institution of afternoon tea in the Computational Neurobiology Laboratory at the Salk, and teatime quickly became the daily occasion for close discussion of ideas and data, flying untried balloons, and giving the broad questions a hearing. It was a time for emerging from the comfortable burrows of safe detail into the wide-open prairie of no-holds-barred. Crick characteristically pushed the questions about how the brain works further and more relentlessly. Moreover, it was typically his hunches, breadth, and steel-edged skepticism that supplied a sense of balance both when we thought we knew what we were doing and when we were pretty sure we didn't. Virtually everyone who visited the Computational Neurobiology Lab was coaxed or bullied into dilating on the philosophical (grand-scale, background, or fuzzy) questions facing computational neuroscience. From these "confessions," we drew ideas and inspiration, and garnered the pluck to stick our necks out a bit.

Several explanations-cum-apologies are in order. The first is for our decision to facilitate easy reading by including only the unavoidable minimum of references in the text itself. We found that long lists of authors in the text make the reader stumble, and hence we elected to use notes for many references rather than follow standard practice in technical writing. Despite our best efforts to refer as fully as possible in the notes, we undoubtedly have missed some essential references, and we apologize in advance for unwitting omissions. The next apology is owed because in choosing instances of research to exemplify a point, we inevitably found ourselves drawing on the research that was most familiar to us, and that often meant research based in California, especially in San Diego. Important and interesting work in computational neuroscience is going on all over the globe, and to have done an exhaustive survey before beginning to write would have meant a deadline receding faster than the progress line. We therefore apologize if we seem rather provincial in our selection preferences. The third apology is for the length. We began the project with the strict understanding that primers are best if brief. In the execution, alas, it became impossible to live within the bounds. As it is, a number of additional topics might well have been included but for permitting the book an embarrassing girth. We therefore apologize—both because the book is too long and because it is too short. Fourth, we decided in the interests of smooth reading to abide by the practice of using "he" as the third-person pronoun referring indifferently to males and females. This reflects nothing ideological. If anything, it is a concession to Mrs. Lundy, whose unflinching dictum in grammar school was that ideological shoe-horning frustrates readability.

Many people helped enormously in writing the book; it simply could not have been done by just the two of us. Most particularly, Paul Churchland gave unstintingly of his imagination and ideas; the daily ritual was to think through everything, page by page, model by model, over capuccino at Il Fornaio. Antonio and Hanna Damasio talked through every major issue with us; they broadened and deepened our perspective in all dimensions, but especially in thinking about what neuropsychological results could tell us about micro-organization. Beatrice Golomb, V. S. Ramachandran, Diane Rogers-Ramachandran, Alexandre Pouget, Karen Dobkins, and Tom Albright helped with representation in general and visual representations in particular; Rodolfo Llinás helped with many issues, but especially in thinking about time; Gyori Buzsáki, Larry Squire, David Amaral, Wendy Suzuki, and Chuck Stevens with plasticity; Carver Mead with thinking about the nature of computation, time, and representation. Shawn Lockery, Steve Lisberger, Tom Anastasio, Al Selverston, Thelma Williams, Larry Jordan, Susan Shefchyk, and James Buchanan gave us much useful advice on sensorimotor coordination. Mark Konishi and Roderick Corriveau gave us invaluable criticism and advice on many chapters and saved us from several embarrassments. Many thanks are also owed to Paul Bush for preparing the glossary, Shona Chatterji for drawing and cheerfully redrawing many figures, Mark Churchland for the cover and for useful criticism, Georg

Schwarz for manuscript preparation, and David Lawrence for rescues from macfrazzles. A special debt is owed to Rosemary Miller, whose wit and wisdom kept the boat afloat. Others who helped in indispensable ways include: Richard Adams, Dana Ballard, Tony Bell, Anne Churchland, Hillary Chase Benedetti, Richard Gregory, Geoff Hinton, Harvey Karten, Christof Koch, Bill Lytton, Steve Nowlan, Leslie Orgel, Hal Pashler, Steve Quartz, Paul Rhodes, Paul Viola, Ning Qian, and Jack Wathey.

P.S.C. was supported by a University of California President's Humanities Fellowship, a grant from the National Science Foundation (87-06757), and the James S. McDonnell Foundation. T.J.S. was supported by the Howard Hughes Medical Institute and grants from the Drown Foundation, the Mathers Foundation, the National Science Foundation, and the Office of Naval Research.

Contents

1 Introduction

Major advances in science often consist in discovering how macroscale phenomena reduce to their microscale constituents. These latter are often counterintuitive conceptually, invisible observationally, and troublesome experimentally. Thus, for example, temperature in a gas turned out to be mean kinetic energy of the constituent molecules; the varied properties displayed by matter turned out to be a function of the component atoms and their arcane properties such as electron shells; bacteria—not Divine vengeance—were found to be the proximal cause of smallpox and bubonic plague; and the reproduction of organisms, we now know, depends on the arrangement of four bases in the molecule DNA.

Our psychological life, too, is a natural phenomenon to be understood. Here as well, the explanations will draw on properties of the infrastructure that are certainly veiled and probably arcane, an infrastructure whose *modus operandi* may seem alien to our customary self-conception. Perhaps this is inevitable, since the very brain we wish to understand is also the brain whose unaided observation is focused at the macrolevel and whose design seems to favor large-scale concepts for the explanation of its own behavior; for example, superstructure concepts such as "is hungry," "wants food," "believes honey is in the hole up the oak tree," and "sees the grizzly bear approaching."

Neurons are the basic structural components of the brain. A neuron is an individual cell, specialized by architectural features that enable fast changes of voltage across its membrane as well as voltage changes in neighboring neurons. Brains are assemblies of just such cells, and while an individual neuron does not see or reason or remember, brains regularly do. How do you get from ion movement across cell membranes to memory or perception in brains? What is the nature of neuron-neuron connectivity and interactivity? What makes a clump of neurons a nervous *system*?

At this stage in the evolution of science, it appears highly probable that psychological processes are in fact processes of the physical brain, not, as Descartes concluded, processes of a nonphysical soul or mind. Since this issue has been discussed at length elsewhere (for example, P. M. Churchland 1984, P. S. Churchland 1986), and since Cartesian dualism is not taken very seriously either in mainstream philosophy or mainstream neuroscience, it is not necessary to repeat the details of the arguments here. Suffice it to say that the

Cartesian hypothesis fails to cohere with current physics, chemistry, evolutionary biology, molecular biology, embryology, immunology, and neuroscience. To be sure, materialism is not an established fact, in the way that the four-base helical structure of DNA, for example, is an established fact. It is possible, therefore, that current evidence notwithstanding, dualism might actually be true. Despite the rather remote possibility that new discoveries will vindicate Descartes, materialism, like Darwinian evolution, is the more probable working hypothesis. That being so, it does not seem worthwhile to modify the basic neuroscience research program and its scaffolding of physicalistic presuppositions to accommodate the Cartesian hypothesis, though scientific tolerance counsels that the door not be closed until the facts themselves well and truly close it. Whether modifications to micro/nano/pico level sciences such as quantum physics will be called for as a result of advances in neuropsychology is likewise conceivable (Penrose 1989), but so far there is no moderately convincing reason to expect that they will.

Arguments from ignorance are to be especially guarded against in this context. Their canonical form is this: neuroscience is ignorant of how to explain X (consciousness, for instance) in terms of the nervous system; therefore it cannot be so explained. Rather, it can eventually be explained in terms of Y (pick your favorite thing, for example, quantum wave packets, psychons, ectoplasmic retrovibrations, etc.). The canonical form lends itself to endless seductive variations, particularly ones in which failures of imagination massage intuition: "We cannot *imagine* how to explain consciousness in terms of neuronal activity . . . ; how could physical processes like ions crossing membranes explain the awfulness of pain?" In its denuded rendition, the argument from ignorance is not mildly tempting, but in full regalia, it may seem beguiling and exactly what reharmonizes such "intuition dissonance" as is provoked by reflecting on the physical basis of the mental. A version of the argument convinced the German mathematician and philosopher, Leibniz (1714), and in the past two decades, variations on Leibniz' basic theme have surfaced as the single most popular and appealing justification for concluding that neurobiological explanations of psychological phenomena are impossible. (For instances of the argument in many different and alluring guises, see Thomas Nagel 1974, J. C. Eccles 1977, John Searle 1980, 1990, and Roger Penrose 1989.) From the revolutions wrought by Copernicus, Galileo, Darwin, and Einstein, it is all too apparent that "intuition dissonance" is a poor indicator of truth; it is a good indicator only of how one idea sits with well-favored others. Establishing truth or probability requires rather more.

The working hypothesis underlying this book is that emergent properties are high-level effects that depend on lower-level phenomena in some systematic way. Turning the hypothesis around to its negative version, it is highly improbable that emergent properties are properties that cannot be explained by low-level properties (Popper 1959), or that they are in some sense irreducible, causally *sui generis*, or as philosophers are wont to say, "nomologically autonomous," meaning, roughly, "not part of the rest of science" (Fodor 1974, Pylyshyn 1984). The trouble with characterizing certain properties as irreducibly emergent is that it assumes we can tell in advance whether something can be explained—*ever* explained. Obviously such a claim embodies a prediction, and as the history of science shows all too clearly, predictions grounded in ignorance rather than knowledge often go awry. In advance of a much more highly developed neurobiology than currently exists, it is much too soon to be sure that psychological phenomena cannot be explained in terms of neurobiological phenomena. Although a given phenomenon such as protein folding or awareness of visual motion cannot *now* be explained, it might yield to explanation as time and science go on. Whether it does or not is a matter of empirical fact, not a matter of *a priori* divination. Searching for reductive explanations of emergent properties does not entail that we should expect the explanations to be simpleminded or breezily cobbled up or straightforwardly readable off the data points; it means only that the betting man keeps going.

Two groundbreaking discoveries in the nineteenth century established the foundations for a science of nervous systems: (1) macro effects displayed by nervous systems depend on individual cells, whose paradigm anatomical structures include both long tails (axons) for sending signals and treelike proliferations (dendrites) for receiving signals (figure 1.1); (2) these cells are essentially electrical devices; their basic business is to receive and transmit signals by causing and responding to electric current. Within this elegantly simple framework, truly spectacular progress has been made in unravelling the intricate story of exactly how neurons work. In this century, and especially within the

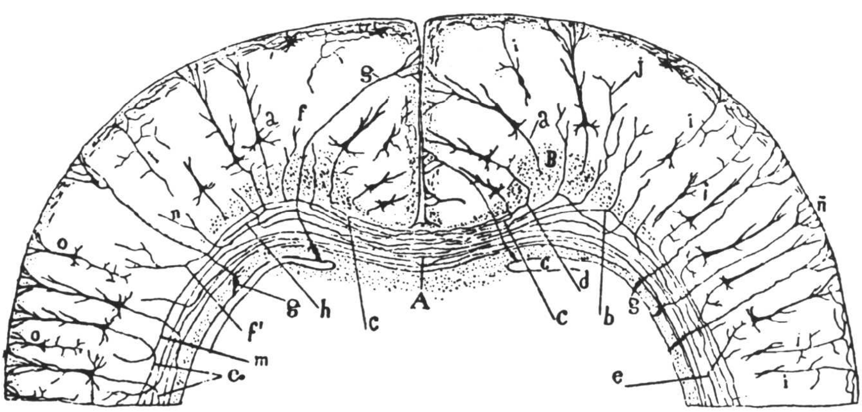

Figure 1.1 Drawing by Cajal based on his Golgi-stained sections of the superior part of the cerebral hemispheres and corpus callosum of a mouse of 20 days. A, corpus callosum; B, anteroposterior fibers; C, lateral ventricle; *a*, large pyramidal cell; *b*, callosal fiber bifurcating into a branch that is arborized in the gray matter and another that continues in the corpus callosum; *c*, callosal fiber that comes from an axon of the white matter; *d*, callosal fiber the originates in a pyramidal cell; *e*, axons of lateral pyramidal cells which follow a descending course in the corpus callosum without forming part of the commissure; *f*, *f'*, the two final branches coming from a fiber of the corpus callosum and arborizing in the gray matter; *g*, epithelial cells; *h*, fiber from a large pyramid, giving off a fine collateral to the corpus callosum; *i*, fusiform cells whose axons ascend to the molecular layer; *j*, terminal arborization of a callosal fiber originating on the opposite side. (With permission. Santiago Ramón y Cajal, 1890. Reprinted in DeFelipe and Jones, eds., 1988, *Cajal on the Cerebral Cortex*. Oxford: Oxford University Press.)

last three decades, an enormous amount has been learned about neurons: about their electrophysiology, microanatomy, connectivity, and development; about the large assortment of neurochemicals that mediate signaling from one neuron to the next; inside a neuron, about the cell's membrane, its roster of channel types, and their specific roles in receiving, integrating, and sending signals; about transmitter release, and about the range, structure, and mechanisms of receptors. Even the genetics of the proteins that constitute the various receptors is now steadily coming into view. (Nathans 1987, 1989, Gasic and Heinemann, 1991, Heinemann et al. 1990).

Recent progress in neuroscience is genuinely breathtaking and deservedly captivating. But, the naif might wonder why, if we know so much about neurons, do we not yet understand how the brain works—or at least how, say, the visual system or the motor system works? Assuming that detailed knowledge of the parts automatically confers (or nearly so) knowledge of the whole, then we ought to understand—more or less, at least in silhouette—how animals see, learn, and take action. In fact, however, we do not. Perhaps the hitch is that microlevel progress notwithstanding, we still do not know nearly enough about the fine-grained neural facts. All that is needed, runs this argument, is more of the same—indeed, much, much more of the same. This strategy is sometimes referred to as the pure bottom-up approach. It counsels that if brains are, after all, just assemblies of cells, then once we truly understand every facet of cell function, the principles of brain function will be evident, by and large. Perhaps. But perhaps not.

The overarching contention of this book is that knowledge of the molecular and cellular levels is essential, but on its own it is not enough, rich and thorough though it be. Complex effects, such as representing visual motion, are the outcome of the dynamics of neural networks. This means that while network properties are dependent on the properties of the neurons in the network, they are nevertheless not identical to cellular properties, nor to *simple* combinations of cellular properties. Interaction of neurons in networks is required for complex effects, but it is dynamical, not a simple wind-up doll affair.

A telling illustration derives from research by Allen Selverston (1988) on the stomatogastric ganglion of the spiny lobster (figure 1.2).[1] The network in question contains about 28 neurons and serves to drive the muscles controlling the teeth of the gastric mill so that food can be ground up for digestion. The output of the network is rhythmic, and hence the muscular action and the grinders' movements are correspondingly rhythmic.

The basic electrophysiological and anatomical features of the neurons have been catalogued, so that the microlevel vitae for each cell in the network is impressively detailed. What is not understood is how the cells interact to constitute a circuit that produces the rhythmic pattern. No one cell is responsible for the network's rhythmic output; no one cell is itself the repository of properties displayed by the network as a whole. Where then does the rhythmicity come from? Very roughly speaking, from the pattern of interactions among cells *and* the intrinsic properties of component cells. What, more precisely speaking, *is* that? How does the network create rhythm? How is it

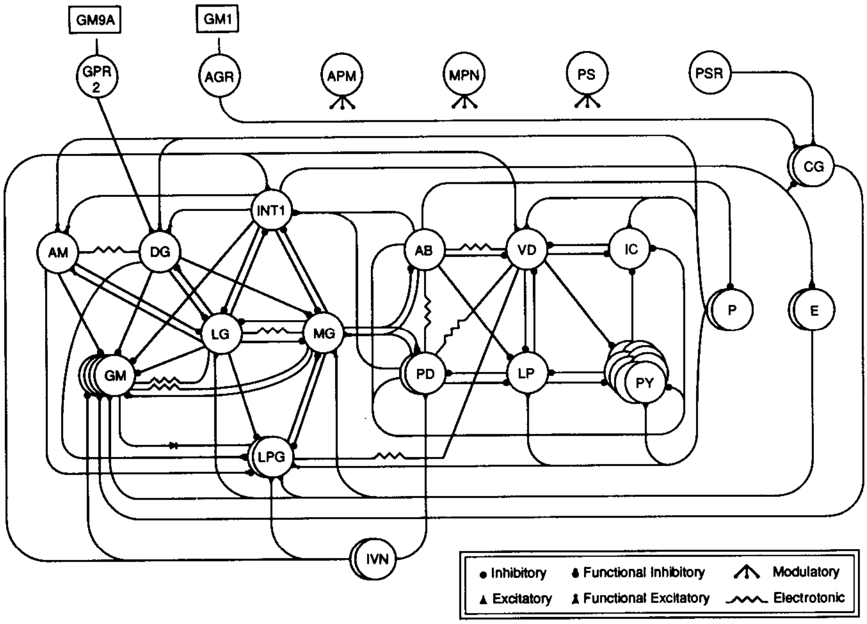

Figure 1.2 Diagram of the circuit in the stomatogastric ganglion of the spiny lobster. The circuit normally has 28 neurons, and for each, its connectivity (whom it affects and who affects it), sign of connectivity (excitatory or inhibitory), and mode of effect (chemical or electrical) have been discovered. Labels on cell bodies stand for their individual names. (Courtesy Allen Selverston.)

that the network can produce different rhythms under different biochemical conditions?

Research on the stomatogastric ganglion is legendary in neurobiology, partly because it is a fair test case for the bottom-up strategy: if the purely bottom-up approach works anywhere, it should work on the stomatogastric ganglion. If the macrolevel answers are supposed to fall out of the microlevel data, they ought to do so here. Yet we are disappointed. As Selverston ruefully points out, the purely bottom-up strategy has all the earmarks of a half-strategy. Moreover, the plea, "If only more microlevel details of the neurons were discovered, then the explanation would be evident," tends now to fall on skeptical ears. What the stomatogastric ganglion seems to be telling us is that we need to figure out the interactive principles governing the system, and that although interactive hypotheses should be constrained by microlevel data, their job is to characterize higher-level features. Boiled down, the lesson is that microlevel data are *necessary* to understand the system, but not *sufficient*. To echo a remark of Maxwell Cowan, even if we did know about all the synapses, all the transmitters, all the channels, all the response patterns for each cell, and so forth, still, we would not know how an animal sees and smells and walks.[2]

There is a broader rationale for modeling that goes beyond neuroscience in particular and applies to science generally. Why bother with models at all, one might ask? Why not just perform experiments and record the observations? Though the answers may be obvious, they are perhaps worth listing. First, models help organize the data and motivate experiments; they suggest how data might fit together to yield an explanation of a phenomenon. It is, therefore, better to have some model than none at all. In fact, of course, scientists do always have some hypothesis or other that provides the motivational and interpretive framework for their research, though background hypotheses may be neither cleanly articulated nor well-honed. A quantitative model is a step forward because it brings background assumptions to the light of day and permits a more exacting analysis of why they might or might not be true. The further philosophical point is that models increase in believability as they survive tough experimental tests (Popper 1959, P. S. Churchland 1986). Especially in the pioneering days of a discipline, when data are relatively sparse, progress is closely tied to *ruling out* a class of models and hypotheses. Indefinitely many models can be equally consistent with a set of data; to make real strides one must seek to falsify an ostensibly plausible model. Consequently, models that suggest potentially falsifying experiments are critical.[3] Should a model survive a demanding experimental test, to that degree it is more probable; saved from the scrap heap of dead hypotheses, it lives on to be tested against yet further experimental data. Should it be falsified, it then becomes a springboard for the next model.

Computational neuroscience is an evolving approach that aims to discover the properties characterizing and the principles governing neurons and networks of neurons. It draws on both neurobiological data and computational ideas to investigate how neural networks can produce complex effects such as stereo vision, learning, and auditory location of sound-emitting objects. To put it crudely, it has one foot in neuroscience and one foot in computer science. A third foot is firmly planted in experimental psychology, and at least a toe is in philosophy, so evidently the enterprise is multipedal. Of which more anon.

Probably the closest academic kin of computational neuroscience is systems neurobiology, a branch of neuroscience that traditionally has focused on much the same set of problems, but did not explicitly ally itself with computer modeling or with an avowedly information-processing framework for theories. A precocious ancestor went by the name of "cybernetics," which, inversely to systems neurobiology, generally leaned more heavily on the engineering and psychophysical sides, and more lightly on the neurobiological side. Coined more recently, "connectionism" usually refers to modeling with networks that bear only superficial similarities to real neural networks, while "neural net modeling" can cover a broad range of projects. Ironically perhaps, "neural net modeling" is usually identified with computer modeling of highly artificial nonneuronal networks, often with mainly technological significance such as medical diagnoses in emergency wards.[4] "PDP" ("parallel distributed processing") is generally the preferred label of cognitive psychologists and some computer scientists who seek to model rather high-level activity such as face recognition and language learning rather than lower-level activity such as visual motion detection or defensive bending in the leech.

As we use the term, "computational neuroscience" aims for biological realism in computational models of neural networks, though *en route*, rather simplified and artificial models may be used to help test and explore computational principles. Academic garden-plotting is a comically imprecise trade because the carrots regularly wander in with turnips and the turnips with the potatoes. Each of us (P.S.C. and T.J.S.) is cheerfully guilty of wandering into neuroscience from his mother discipline, so we emphatically do not mean to tut-tut academic "cross-fielding." On the contrary, we view the blurring of the disciplinary boundaries between neuroscience, computer science, and psychology as a healthy development to be wisely encouraged. In any case, perhaps a crude survey will help orient the greenhorn—or even the old hand—to the clustering of goals, tactics, and prejudices manifest in the "network" game.

The expression "computational" in computational neuroscience reflects the role of the computer as a research tool in modeling complex systems such as networks, ganglia, and brains. Using the word in that sense, one could have also computational astronomy or computational geology. In the present context, however, the word's primary force is its descriptive connotation, which here betokens the deep-seated conviction that what is being modeled by a computer is itself a kind of computer, albeit one quite unlike the serial, digital machines on which computer science cut its teeth. That is, nervous systems and probably parts of nervous systems are themselves naturally evolved computers—organically constituted, analog in representation, and parallel in their processing architecture. They represent features and relations in the world and they enable an animal to adapt to its circumstances. They are a breed of computer whose *modus operandi* still elude us but are the mother lode, so to speak, of computational neuroscience.

A number of broad clues about computation in nervous systems are available. First, unlike a digital computer which is general purpose and can be programmed to run any algorithm, the brain appears to be an interconnected collection of special-purpose systems that are very efficient at performing their tasks but limited in their flexibility. Visual cortex, for example, does not appear able to assume the functions of the cerebellum or the hippocampus. Presumably this is not because visual cortex contains cells that are essentially and intrinsically visual in what they do (or contain "visons" instead of "auditons"), but rather it is mainly because of their morphological specialization and of their place in the system of cells in visual cortex, i.e., relative to their input cells, their intracortical and subcortical connections, their output cells, and so on. Put another way, a neuron's specialization is a function of the neuron's computational roles in the system, and evolution has refined the cells better to perform those roles.

Second, the clues about the brain's computational principles that can be gleaned from studying its microstructure and organization are indispensable to figuring out its computational organization because the nervous system is a product of evolution, not engineering design. Evolutionary modifications are always made within the context of an organization and architecture that are already in place. Quite simply, Nature is not an intelligent engineer. It cannot dismantle the existing configuration and start from scratch with a preferred design or preferred materials. It cannot mull the environmental conditions and construct an optimal device. Consequently, the computational solutions evolved by Nature may be quite unlike those that an intelligent human would invent, and they may well be neither optimal nor predictable from orthodox engineering assumptions.

Third, human nervous systems are by no means exclusively cognitive devices, though the infatuation with cognition fosters a tacit tendency to assume so. Nervous systems must also manage such matters as thermoregulation—a very complex function for mammals—growth, aspects of reproduction, respiration, regulation of hunger, thirst, and motor control, and maintenance of behavioral state, such as sleeping, dreaming, being awake, and so forth. Thus an evolutionary modification that results in a computational improvement in vision, say, might seem to have the earmarks of an engineering prizewinner. But if it cannot mesh with the rest of the brain's organization, or if it marginalizes critical functions such as thermoregulation, the animal and its "prize-winning" vision genes will die. Given these reasons, *reverse* engineering, where the device is taken apart to see how it works, is a profitable strategy with respect to the brain. By contrast, a purely *a priori* approach, based entirely on reasonable principles of engineering design, may lead us down a blind alley.

Fourth, it is prudent to be aware that our favorite intuitions about these matters may be misleading, however "self-evident" and compelling they be. More specifically, neither the nature of the computational problems the nervous system is solving nor the difficulty of the problems confronting the nervous system can be judged merely by introspection. Consider, for example, a natural human activity such as walking—a skill that is typically mastered in the first year or so of life. One might doubt whether this is a computational problem at all, or if it is, whether it is a problem of sufficient complexity to be worth one's reflection. Since walking is virtually effortless, unlike, say, doing algebra, which many people do find a strain, one might conclude from casual observation that walking is a computationally easy task—easier, at least, than doing algebra. The preconception that walking is computationally rather trivial is, however, merely an illusion. It is easy enough for toy manufacturers to make a doll that puts one foot in front of the other as long as she is held by the child. But for the doll to walk as we do, maintaining balance as we do, is a completely different task. Locomotion turns out to be a complicated matter, the ease implied by introspection notwithstanding.

Another computational issue of critical importance in generating hypotheses in computational neuroscience concerns the time available for performing the computation. From the point of view of the nervous system, it is not enough to come up with solutions that merely give the correct output for a given input. The solutions must also be available within milliseconds of the problem's presentation, and applications must be forthcoming within a few hundred milliseconds. It is important that nervous systems can routinely detect signals, recognize patterns, and assemble responses within one second. The ability of nervous systems to move their encasing bodies appropriately and swiftly was typically selected at every stage of evolution, since by and large natural selection would favor those organisms that could flee or fight preda-