A New Era in Computation

edited by N. Metropolis and Gian-Carlo Rota

The MIT Press
Cambridge, Massachusetts
London, England
1993

Terrence J. Sejnowski and Patricia S. Churchland

Computation in the Age of Neuroscience

BRAINS ARE MASSIVELY PARALLEL, ANALOG, biological computers. They are far more powerful, flexible, and compact than any manufactured computer. The computational principles that make brains such effective computers are radically different from those that are used in conventional digital designs.[1] Some neurocomputational principles, such as analog processing in dendritic trees and synaptic plasticity, have already been discovered by biologists, but many more principles have yet to be identified. Analog VLSI (very large scale integration) technology may provide a medium for exploring these principles and creating new computational architectures. A partnership between neuroscience and computer technology is opening a path toward building silicon brains that can grapple with the real world. In the next century, this partnership could provide new insights into the nature of our own brains as well as engender ideas for remarkable new machines made in our image.

BEES AND AGRIBOTS

Imagine a computing device that would revolutionize the stoop labor sector of agriculture. In midsummer, the agribot travels up and down rows of tomato plants, picking those tomatoes reddened to maturity while leaving small green tomatoes on the vine. It sorts the harvest according to size, tossing down for fertilizer

---

Terrence J. Sejnowski is an Investigator with the Howard Hughes Medical Institute, Director of the Computational Neurobiology Laboratory at the Salk Institute in La Jolla, and a member of the Biology Department at the University of California, San Diego.

Patricia S. Churchland is a MacArthur Fellow in the Philosophy Department at the University of California, San Diego and an adjunct member of the Salk Institute.

any fruit spoiled by rotten spots. Its sister agribots rove the tomato fields in early spring, some pulling up weeds and some picking off grubs, thereby reducing the farm's dependence on pesticides and herbicides. Other agribots toil in the peach orchards; some prune the trees in early spring, some delicately pick the fruit during the summer.

Although mechanical devices, including tomato-pickers, have been invented for agriculture, typically the machine-product interface requires downgrading the produce so that nondiscerning and inflexible machines can be used. But this compromise is not inevitable. Unafraid of heights and unwearying in their chores, the agribots allow farmers to grow high quality, as opposed to rubbery and tough, fruit and vegetables.

What would it take to make this appealing science fiction real? First, the agribots need computational insides that are *very* small, *very* cheap, and *enormously* powerful. Second, the computational style of the devices must allow for flexibility in sensory categorization and adaptability in motor control. Abandoning the conventional wisdom—"program a universal machine to create a precise virtual machine"— engineers will have to harken to a very different strategy: "create a real-time machine with feedback mechanisms, and let it learn to perform its task." Is this fantasy within the realm of engineering possibility?

Superficially, the bets would appear to be against. Existing digital computers lack the autonomy, flexibility, and adaptability required by the fictional agribots. Nothing remotely close to the desired miniaturization and energy efficiency is available in the electronics marketplace.

A more positive response, however, comes from noting that biological computers replete with the agribot's features do already exist. Brains of insects, birds, fish, and mammals represent an existence proof for powerful, fast, flexible, and self-reliant computers. Nature did it, so it ought to be possible for us to do it. Nevertheless, between the existence proof and the construction kit lies a vast gap.

Consider, for example, the brain of a honeybee, which has about a million neurons compared with the 100 billion neurons in a human brain (figure 1). Consider *energy efficiency*: the bee's brain dissipates less than 10 microwatts (10^{-6}), superior by about

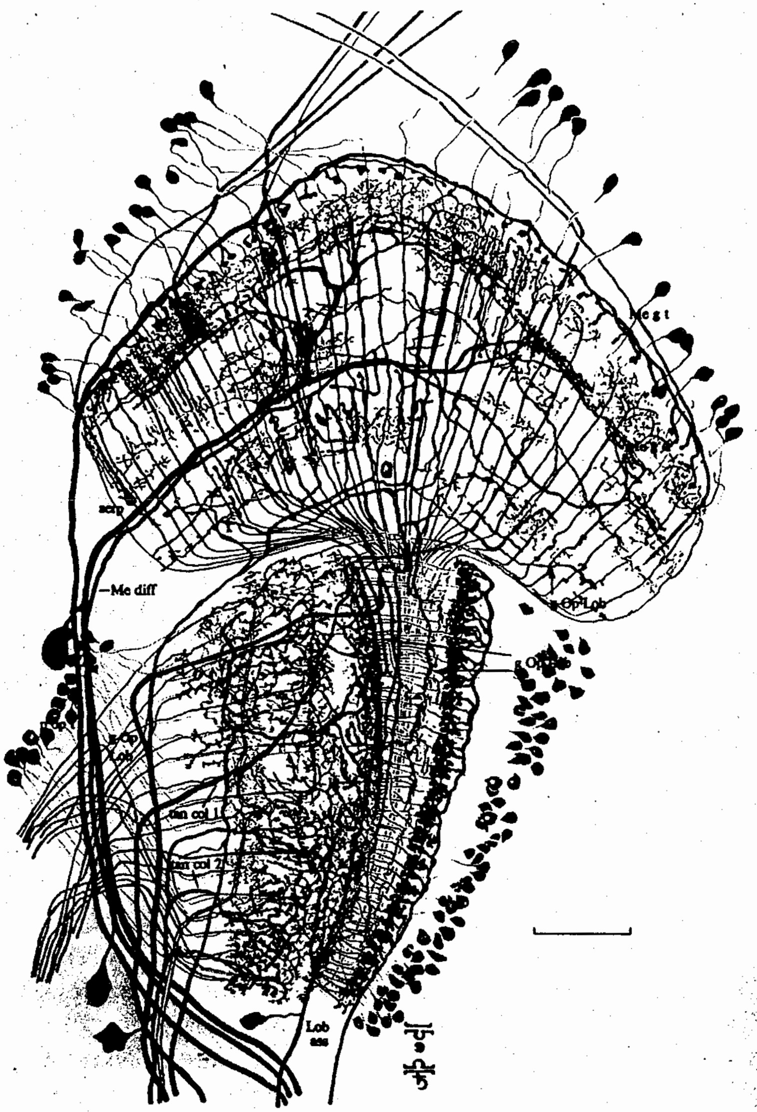

Figure 1. Circuitry in the visual system of the fly (*Musca domesticus*). The organization is very like that in bees, butterflies, and other insects. This drawing shows only about 10% of the actual number of neurons in the region. Even without identifying each pathway and each neuron type, the system's highly organized and regular structure is visible. Note that the large neurons (black, left) differ systematically in the level of their destination, the choice of cells to contact, and the precise location of contact with specific cells. Scale bar is 20 µm. (Drawing by N. J. Strausfeld, *Atlas of an Insect Brain*. New York: Springer-Verlag, 1976)

7 orders of magnitude to the most efficient manufactured computer. Consider *speed*: a bee brain, on a rough and conservative estimate, performs about 10 teraflops (a thousand gigaflops); the most powerful computers approach speeds of only 10 gigaflops (a billion operations per second). Consider *behavioral repertoire*: honeybees harvest nectar from flowers and bring the nectar back to the hive. They maximize foraging benefits and minimize foraging costs, for example, by recognizing high nectar sites and remembering which flowers have been visited. They can see, smell, fly, walk, and maintain balance. They can navigate long distances, and they can predict changes in nectar location by extrapolating from the past regularities. They communicate nectar sources to worker bees in the hive, they recognize intruders and attack, they remove garbage and dead bees from the hive, and when the hive becomes crowded, a subpopulation will swarm in search for a new home. Consider *autonomy and self-reliance*: bees manage all this on their own, without the aid of human intelligence to supervise them, repair them, or nurse them along. A supercomputer, by contrast, needs the constant tender care of a cadre of maintainers and programmers, not to mention financiers. Finally, the entire bee brain takes up only about a few cubic millimeters of space, a marvel of miniaturization.

Research aimed at understanding how brains work will likely be a profitable undertaking, even in the near future. On the one hand are important medical consequences, not to mention the sheer intellectual value of understanding ourselves. Less appreciated, but potentially more significant in economic terms, are technological spin-offs. Knowledge of evolution's computational tricks and architectural ingenuities for speed, power, and flexibility can be applied to a variety of problem domains.

Agriculture may be but one such domain; deep sea mining, defense, pollution clean-ups, and space exploration are others. Current generation neural nets capture some highly general features of brains, such as parallel architecture. Primitive as they are, neural nets of 1990 vintage might be epoched as the "Bronze Age" artifacts of brain-style computer technology. This is not to belittle their very significant beginning, but only to affirm that they are *just a beginning*. Still ahead are epochs advancing the art beyond this first step.

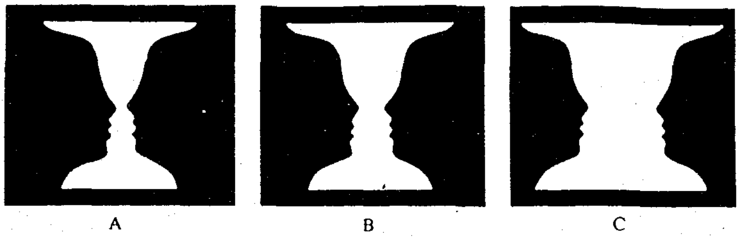

Figure 2. Figure-ground reversal. There are two perceptual interpretations of these images: a pair of black faces, or a white vase. The interpretation can be influenced by conscious attention and biased by features in the image. Thus, the face interpretation is generally favored in A and the vase interpretation in C. (From S. Coren and L. W. Ward, *Sensation and Perception*. San Diego: Harcourt Brace Janovich, 1989)

REVERSE ENGINEERING THE BRAIN

In a very general sense, brains compute. The net result is that they represent the world and manage to survive in it. Computational neuroscience is an emerging field dedicated to figuring out how real brains represent and compute. Computer modeling of neural circuits is essential to the enterprise as a means of addressing how neurons (the cellular components of nervous systems) interact with each other to produce complex effects such as segregating figure from ground, recognizing a banana in many different orientations, or visual tracking of moving targets in 3-D space (figure 2).

Neuroscience contributes three main ingredients to the neuromodeling enterprise: (1) anatomical parameters such as the precise tree structure of various neuron types and the exact connectivity between neurons (who talks to whom) in a particular real network; (2) physiological parameters such as response characteristics of neurons, time constants, synaptic sign, etc., and (3) clues to the network's function and its computational *modus operandi* in executing that function. The main techniques involve intervening in the system, for example by lesioning or by electrical stimulation (figure 3).

Models highly constrained by neurobiological parameters provide a particularly efficient means for exploring the computational wherewithal of nervous systems. Analysis of a highly constrained working model, in turn, can inform neurobiologists about unsus-

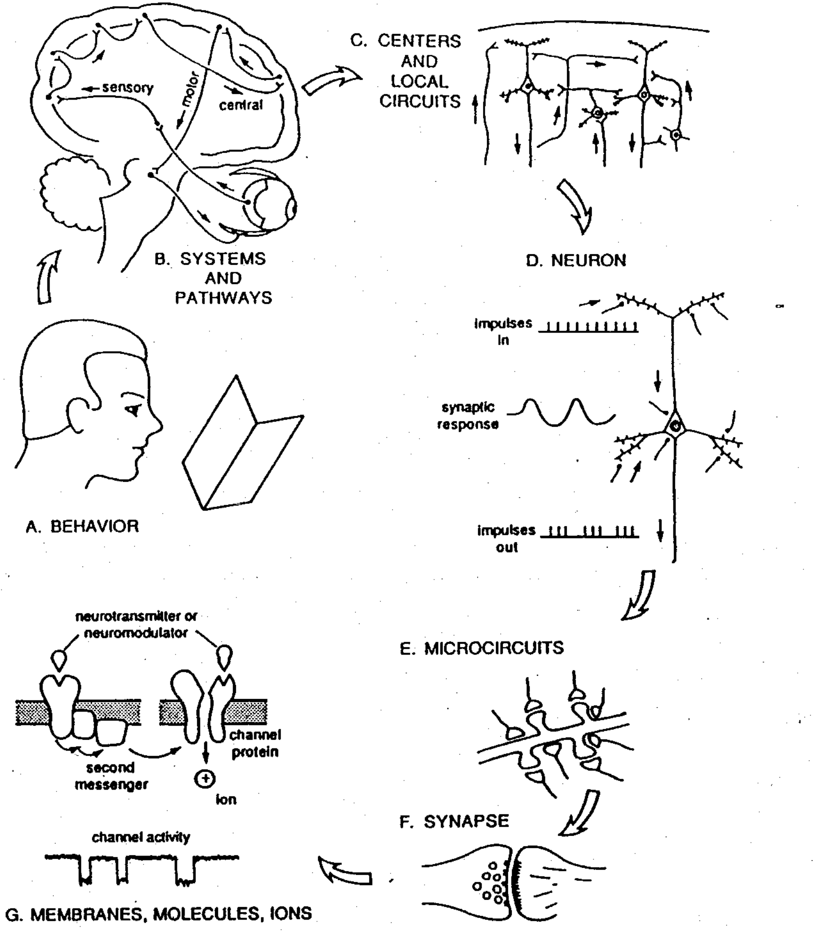

Figure 3. Levels of organization in nervous systems. According to spatial scale, the components of the nervous systems as currently understood include (A) *central nervous system behavior,* including the brain and spinal cord; (B) *systems,* such as the visual system, auditory system, and motor system; in many regions of the brain the topography of neurons corresponds to the topography of its input domain, such as the retina or the skin; (C) *networks* of neurons, which may consist of many thousands of interconnected neurons; (D) the individual *neuron,* which can be either excited or inhibited by inputs from other neurons. (E) On small patches of dendrite of a given cell, the input to synapses interact to constitute a *microcircuit.* (F) At the *synapse,* a signal is passed from the sending cell to the receiving cell, usually by means of a chemical released from the sender that attaches to the receiving cell and changes the cell's voltage a tiny amount. (G) On the *molecular* level, ion channels consist of proteins in the membrane that may reconfigure under restricted conditions to allow specific ions, for example Ca^{2+}, to enter the cell in response to chemical or electrical signals. (From G. Shepherd, Neurobiology, 2nd edition. Oxford: Oxford University Press, 1988)

pected mechanisms and interactions, whose reality can then be tested in actual nervous systems.

By continuously inspiring, correcting, and informing each other, neuroscience and computer modeling can co-evolve to greater accuracy and greater adequacy, respectively. This co-evolution is already producing ideas for innovative computational procedures and architectural design relevant to such problems as real-time interacting, efficient associative memory storage, mixed modality coordination, multiplexing, and attentional selectivity.

SIMULATE OR SYNTHESIZE?

As in simulating a hurricane or a heart on a digital machine, simulating neurons runs afoul of the real time problem. Such machines are not yet powerful enough to both faithfully simulate the system's processes and do it in real time. Either close imitation of neuronal operations or real time has to be sacrificed. The problem is that the simulation strategy consists of compartmentalizing the phenomenon and solving vast numbers of differential equations. It is, therefore, pitifully slow, relative to the real performance-time of the system simulated.

In a neuron, ions pass back and forth across the membrane, signals are integrated, output spikes are produced[2]—all in a matter of a few milliseconds (figure 4). To simulate a millisecond in the life of a neuron, however, thousands of coupled nonlinear differential equations have to be solved. To compound the difficulty, these equation are "stiff," in the sense that they embrace a wide variety of time scales. This means that the simulation's time steps can only be as long as the *shortest* significant interval. Consequently, even a powerful work station will take minutes to simulate a millisecond of real time of a single neuron (figure 5).

Constructing dedicated hardware for synthetic neurons and synthetic nervous systems is the way to circumvent this dead-ending ponderousness of simulation. One obvious strategy, then, is to construct neuron-like chips. To construct chips that compute as wondrously as real neurons, we must first understand how real neurons do it. The production of a spike in the axon of a neuron is indeed an all-or-nothing affair, but the purely digital properties of neurons stop there. Processing in dendrites is analog, including

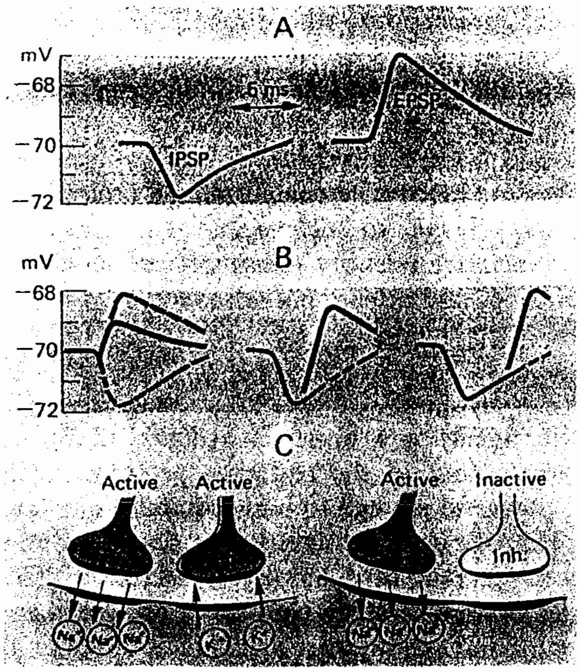

Figure 4. Inhibitory and excitatory synapses on a neuron. (A) The inhibitory postsynaptic potential (IPSP) means that the postsynaptic cell hyperpolarizes (dropping from –70 mv to –72 mv), and the excitatory postsynaptic potential (EPSP) means that the postsynaptic cells depolarizes (from –70 mv to –67 mv). (B) The EPSP was triggered about 1, 3, and 5 sec after the onset of the IPSP. (C) The subsynaptic conductance changes occurring when excitatory and inhibitory synapses are activated simultaneously (left) and when only the excitatory synapse is activated (right). (From R. F. Schmidt, *Fundamentals of Neurophysiology.* Berlin: Springer-Verlag, 1978)

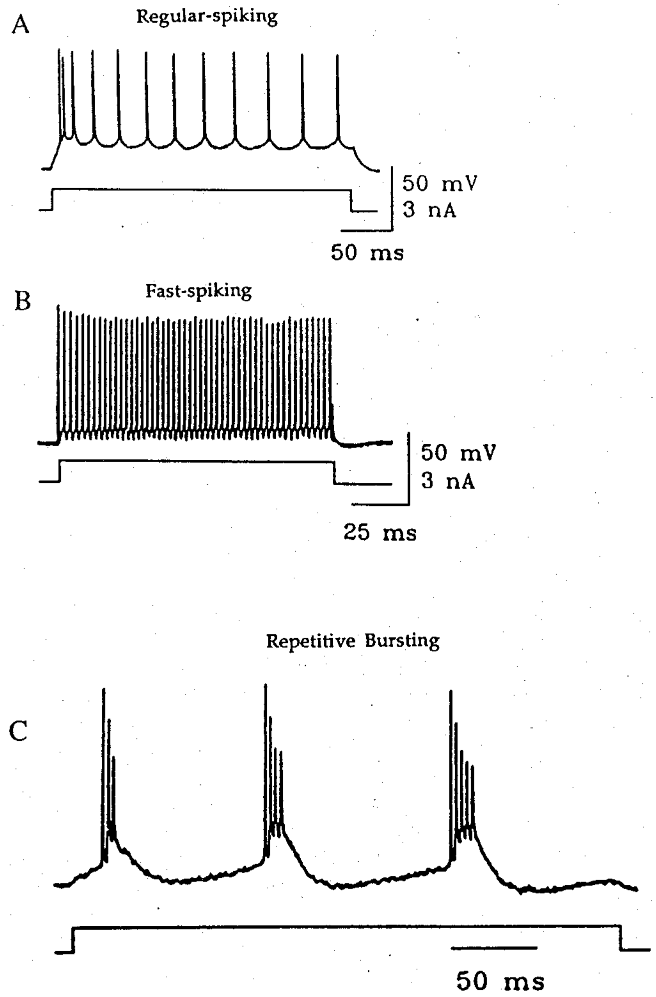

Figure 5. Differences in intrinsic firing patterns of cortical neurons. (A) When stimulated with a suprathreshold step of depolarizing current, regular-spiking neurons respond with an initial high-frequency spike output that rapidly declines to much lower sustained frequencies. Intracellular voltages are displayed in the top trace, injected current steps in the bottom trace. (B) Under similar conditions, fast-spiking cells generate high frequencies that are sustained for the duration of the stimulus. (C) Repetitive intrinsic bursting to a prolonged stimulus. Mean inter-burst frequency was about 9 Hz. (From B. W. Connors and M. J. Gutnick, "Intrinsic firing patterns of diverse neocortical neurons." *Trends in Neurosciences* 13, 1990: 98–99).

both active nonlinear mechanisms that amplify signals as well as passive cable properties. Even axonal spiking is analog in some respects; the time when the spike occurs (it can be any time), the frequency of spikes in a train, the duration of the repolarization period, and so on (figure 6).

Real circuits, of course, have many imperfections. Invariably, they lapse from idealizations and component homogeneity, membranes are typically leaky, components malfunction or drop dead, there can be cross-coupling, and so forth. Contrary to the impulse to shun chip construction in favor of simulation, the farseeing advice is that we find the secret of how to get precision, speed, and power out of imperfect and imprecise components. The point is, somehow neurons operate in real time and cope magnificently, probably *exploiting* "imperfections" to their advantage. Consequently, the coping capacity of real-world neurons is itself computationally interesting.

Neurons are organic. They use lipid molecules to make resistive membranes, complex proteins for ion channels that allow current flow across the membrane, and cytoplasm in which current travels; mitochondria are the micro powerpacks, circulating oxygen is the energy source, and so on. What can the engineer use to construct synthetic neurons, if not these materials?

THE TECHNOLOGY OF CHOICE: ANALOG VLSI

Analog VLSI technology turns out to be well suited to the construction task for two reasons, one theoretical and one practical.[3] (1) The device physics of doped silicon, when operating in subthreshold regions, is comparable to the biophysics of ion channels in neuron membrane; that is, the current passing through a membrane's ion channels follows Boltzmann statistics. This allows one to implement the differential equations directly with analog circuits in CMOS (complementary oxide semiconductor) VLSI. (2) The very same techniques used for creating digital VLSI chips can be adapted to make analog VLSI chips. Carver Mead (at Cal Tech and Synaptics) and Federico Faggin of Synaptics, who played leading roles in digital chip technology, are now spearheading the development of analog chip technology for neural systems (table 1).

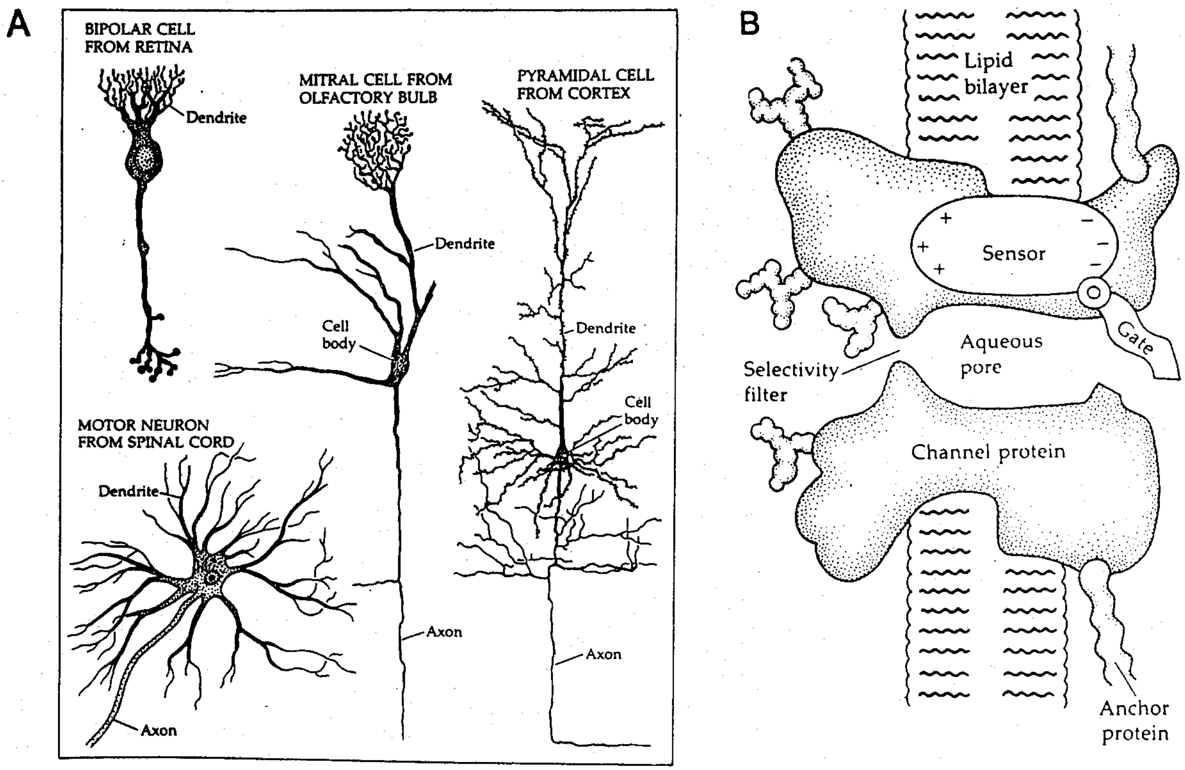

Figure 6. Neurons and neural mechanisms. (A) Examples of neurons illustrating the variety of shapes in different areas of the brain. (From S. W. Kuffler, J. G. Nicholls and A. R. Martin, *From Neuron to Brain: A Cellular Approach to the Function of the Nervous System,* 2nd edition. Sunderland, Mass.: Sinauer, 1984) (B) Working hypothesis for a voltage-gated channel. The transmembrane protein is shown with a pore that allows sodium ions to flow between the extracellular and intracellular sides of the membrane when the gate is open. (From B. Hille, *Ionic Channels of Excitable Membranes.* Sunderland, Mass.: Sinauer, 1984)

TABLE I. VLSI neural nets: figures of merit. Analog VLSI is strikingly superior to digital technology in terms of cost, power and computation density. (From Federico Faggin)

	Cost (MCS[a]/\$)		Power (MCS/Watt)		Computation Density (MCS/ft^3)	
	1991	2000	1991	2000	1991	2000
Conventional digital	.002	.1	.1	10	.2	10
Special purpose digital	.1	4	10	10K	10	1000
Dedicated digital	5	200	500	50K	40	3K
Dedicated analog	500	20K	50K	5M	4K	4M
Human brain[b]		10^9		10^{10}		10^{11}

[a] MCS = million connection updates/second
[b] Assumes that the cost of a human brain is \$10M

Digital technology is still very much in its heyday. It dominates not only the marketplace but also the imagination-space most people explore in thinking about a problem. The tremendous potential of analog VLSI, especially in addressing messy real-world problems, as opposed to made-exact bench problems, has yet to be fathomed. For example, current algorithms running on a digital machine can perform at about 60% reading written numerals on credit card sales forms. They do this well only because the sales slip "exactifies" the data: numerals must be written in the blue boxes. This serves to segment the several numerals, guarantee an exact location, and very narrowly limit numeral size. The blue boxes mean that the really trenchant problem of segmentation does not have to be solved by the machine.

By contrast, machine reading zip codes on mail is an essentially unsolved problem even in the relatively tidy case where the numerals are machine-printed rather than handwritten. The trouble is, the preprocessing regimentations for numeral entry on the sales slips largely do not exist in the mail-world. Here the computer readers have to face the localization problem (where are the numerals and in what order?) and the segmentation problem (what does a squiggle belong to?) as well as the recognition problem (is it a 0 or a 6?).

The crux of the difficulty is that digital machines typically serialize the problem, so naturally they are programmed first to solve the segmentation problem and after that, to solve the recognition problem by a template procedure. Should the machine missolve or fail to solve the segmentation problem, recognition is doomed. In the absence of strict standardization of location, font, size, relation to other numerals, relation of zip code to other lines, and so forth, digital machines regularly fumble the segmentation problem. The best can not yet read even 50% of real mail presented.

Brains, it appears, do not serialize the segmentation and recognition problems in lockstep fashion. Often as not, recognitional cues are used to solve the segmentation problem. In general the brain's approach looks more like cooperative computation or constraint satisfaction than like theorem-proving. Interactive problem solving appears to be typical in the nervous system, whether the problem is speech recognition, object recognition, or organizing a bodily movement during prey-catching. A major advantage of analog VLSI is that the chip can follow the brain's example, solving the segmentation and recognition problems concurrently.

SILICON NEURONS

The first step in building silicon neurons was reported by Mahowald and Douglas, in *Nature* 1992. Using analog VLSI, they created a chip with selected prominent properties of pyramidal neurons, a type found in cortical structures whose basic properties have been intensively studied by neuroscientists. Their silicon neuron was highly simplified, consisting of only one compartment (an axon) and four types of ion channels in the axon membrane. A real pyramidal neuron, by contrast, might have thousands of dendritic segments as well as an axon, tens of thousands of synapses, and scores of types of ion channels (figure 7).

As a pilot project, however, it rates as successful on several counts. First, it ran in real time. A bonus of achieving the difficult goal of real-timeliness is that Mahowald and Douglas could also

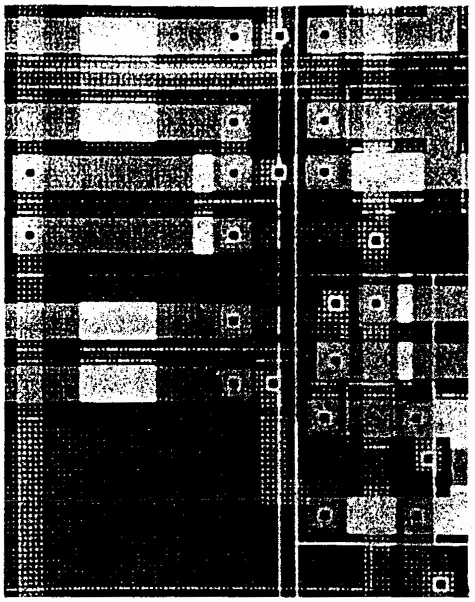

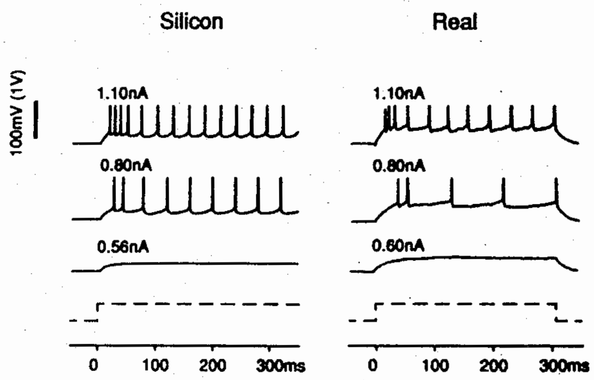
Silicon
Real
100mV (1V)
1.10nA
0.80nA
0.56nA
1.10nA
0.80nA
0.60nA
0
100
200
300ms
0
100
200
300ms

Figure 7. Comparison between the output from a silicon neuron and a biological neuron. (Top left) Magnified view of a part of a VLSI circuit for a silicon neuron. (Top right) Cortical neuron stained with the Goldi technique and a recording microelectrode. (From D. Hubel, "The brain," *Scientific American* Sept. 1979: 44–53). (Bottom left) Response of the silicon neuron to depolarizing current steps. A subthreshold current step (broken lines) injected into the cell body compartment of the silicon neuron evokes a passive charging curve. Larger inputs evoke adapting discharges of action potentials. Response of a cortical pyramidal cell recorded in vitro is provided for comparison. (From M. Mahowald and R. Douglas, "A silicon neuron." *Nature* 354, 1991: 515–518.)

conduct experiments by tweaking many parameters in real time, such as density of channels of a given type, or kinetic rate constants for channel opening. Second, the silicon neuron's output behavior for various injections of current, as displayed by traces on an oscilloscope, closely resembled the output of a real pyramidal cell under various physiological conditions. Third, the Mahowald-Douglas neuron consumes very little power.

With the successful debut of the single, simplified synthetic neuron, a number of developments are on the agenda. One is to upgrade the inaugural realism, for example by adding more compartments (corresponding to dendrites), and a wider range of ion channels. Another step is to build many neurons on a single chip. These synthetic neural circuits can then be explored to learn more about the computational possibilities inherent in the interaction of the various parameters.

A major challenge on this front is the interface between the human creator and the chip. Ideally, the scientist should be able to tweak thousands of parameters in real time, and hence the interface has to be flexible and user-friendly. Using synthetic circuits as experimental preparations means researchers can explore virtual neural-reality rather than sit at a work station watching points appear on a graph.

A further refinement is to make the chip modifiable by "experience" so that hand-setting of neuronal connectivity can be replaced by a training regimen. Mead and his group are currently developing trainable chips, where connectivity is modifiable according to learning rules similar to those believed to underlie plasticity in nervous systems, such as the Hebb rule. Ultimately, one will want to create chips with subpopulations of neurons specialized for different tasks, in the manner that distinct brain regions—visual cortex, auditory cortex, motor cortex, etc.—are specialized. Here again, fruitful ideas may come from seeing how Nature engineers specialization and integration.

Following Nature may require that we model patterns of neuronal connectivity, both long range (on the order of centimeters) and short range (millimeters). Nervous systems are remarkably fault tolerant, in the sense that a circuit and its function can survive quite well the death of individual neurons in the circuit. Comparable fault tolerance might be achievable for artificial systems by imitating the connectivity, modifiability, and processing style of the brain.

NEURAL CIRCUITS IN SILICON

Peripheral sensory organs such as the eye are highly specialized for transduction of external physical signals, such as photons, into electricity activity. The retina is a powerful preprocessor that transforms information about photons into a form suitable for neural representation and computation.[4] In many animals, sensory transducers and preprocessors have evolved to a sensitivity close to the limits of physical possibility. In primates, photoreceptors in the retina will respond to a few photons; the human ear is close to the limits set by the noise from Brownian motion. These transducers do not discretize time as digital computers must do but use powerful analog preprocessors to shape the information to a neural-friendly form. Can these inventions of nature be reverse engineered?

Carver Mead has built a family of silicon retinae. Each is a VLSI chip, merely a square centimeter in area; it weighs about a gram, and it consumes only about a milliwatt of power. Between arrays of phototransistors etched in silicon, dedicated circuits execute smoothing, contrast enhancement, and motion processing. The chip operates in its subthreshold, analog mode. Compared with a typical CCD (charge-coupled device) camera and standard digital image processor, the Mead chip is a paragon of efficiency, power, and compactness. The special-purpose digital equivalent would be about as large as a washing machine. Unlike cameras that must time-sample, typically 60 frames per second, the analog retina works continuously without the need to sample until the information leaves the chip already preprocessed (figure 8).

The operations performed on the current generation of chips capture some of the operations performed by real retinae. For example, resistive grids mimic the function of the layer of horizontal cells that provide lateral interaction between photoreceptors to effect smoothing. There are, however, many more circuits in real retinae that are not included in Mead's synthetic retina, such as amacrine cell circuits. About 30 types of amacrine cells exist in the retina, some are known to perform temporal filtering

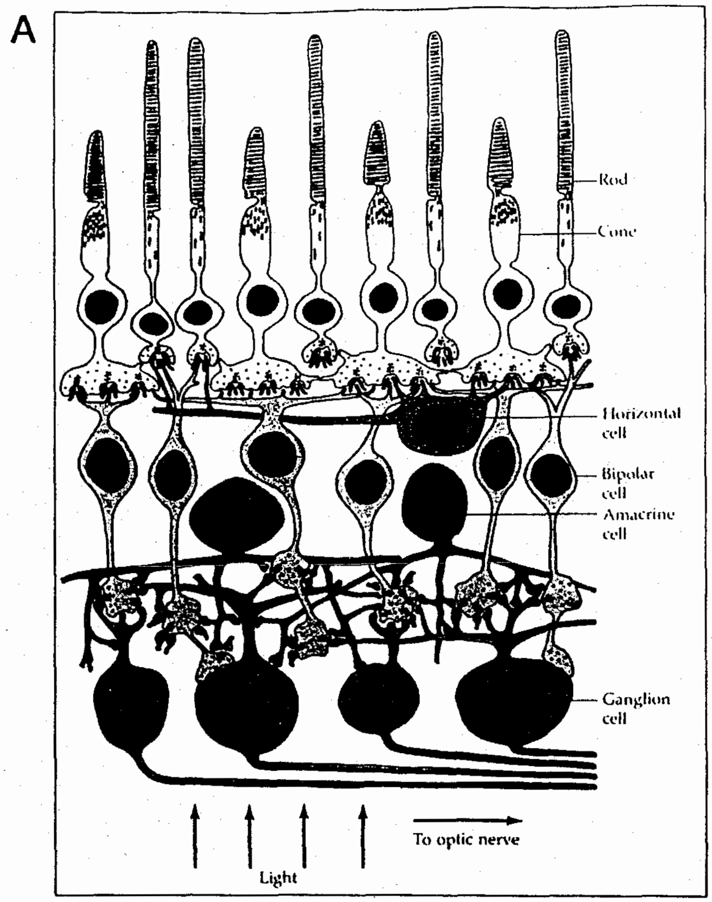

Figure 8. Comparison between the organization of biological and silicon retinae. (A) Diagram showing a close-up of a tiny region on the retina that illustrates several prominent cell types. The outer plexiform layer contains synaptic connections between photoreceptors, horizontal cells that provide lateral interactions, and bipolar cells that carry signals to the ganglion cells. (B) Horizontal cells in the white perch retina. (Left) A single horizontal cell injected with a fluoresecent dye. (Right) Horizontal cells have reciprocal connections as revealed here by dye coupling from one cell. (From J. E. Dowling [1987] *The Retina: An Approachable Part of the Brain.* Cambridge: Harvard University Press.) (C) Diagram of the silicon retina showing the resistive network similar in its function to the array of horizontal cells in the retina; a single pixel element is illustrated in the circular window. The silicon model of the triad synapse consists of a follower-connected transconductance amplifier by which the photoreceptor drives the resistive network, and an amplifier that takes the difference between the photoreceptor output and the voltage stored on the capacitance of the resistive network. These pixels are tiled in an hexagonal way. The resistive network results from an hexagonal tiling of pixels. (From C. Mead [1989]. *Analog VLSI and Neural Systems.* Reading, Mass.: Addison-Wesley.)

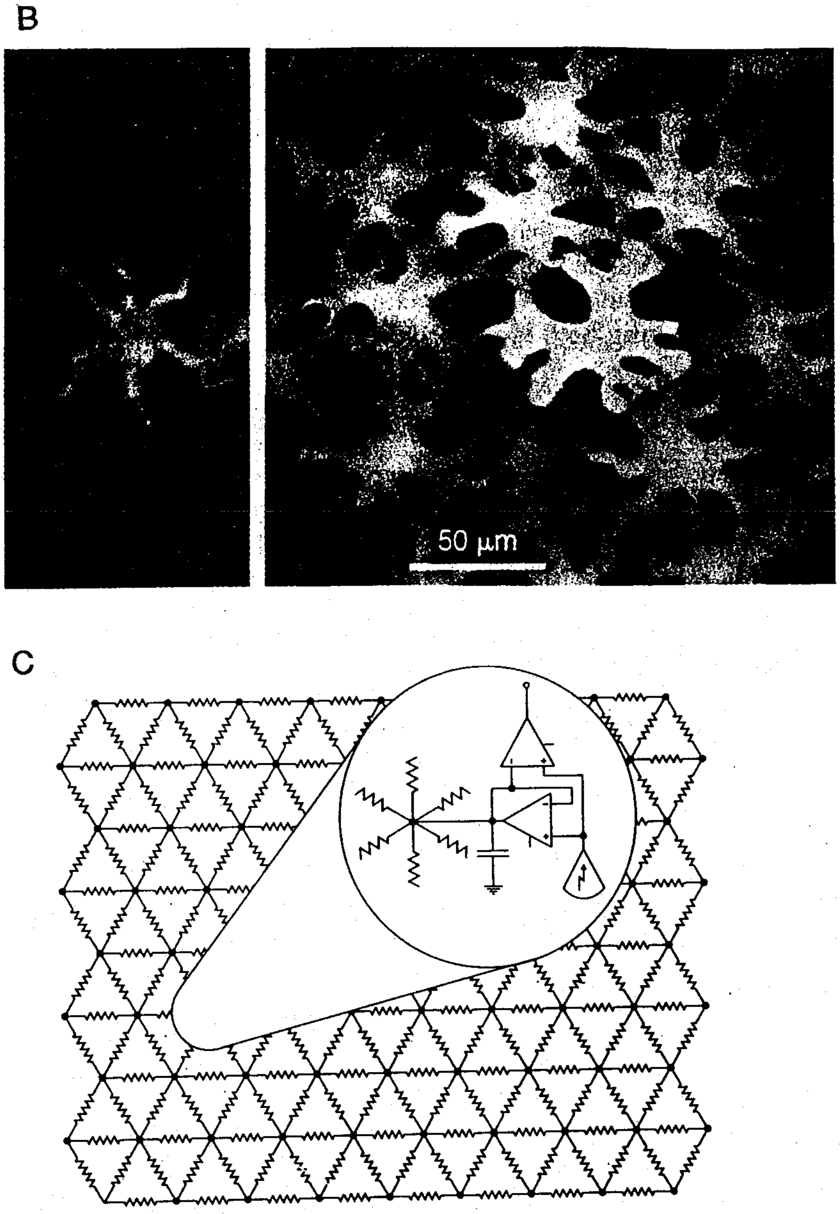
B
50 µm
C

and to provide adaptational mechanisms. So far, however, the precise function of most amacrine cells is not known. The retina is a prime place for the co-evolution of chip design and experimental neurobiology.

To maximize efficiency in the three critical dimensions (power, cost, and density), it makes good sense to build chips. To approximate neural efficiency, however, the technology still has a *very* long way to go. The commercial incentive to push forward with analog VLSI technology will depend on whether the commercial payoff in the long haul looks promising.

SENSORIMOTOR CONTROL

Visual perception in animals, as we all know, is breathtakingly difficult computationally. So far nothing in computer vision has come close even to matching the visual capacity of a bee, let alone that of a rat. In addition to good visual perception, robots will require good coordination between their "eyes" and their "hands." Several observations on the neurobiology of sensorimotor control may provoke new computational insights.

Three features of neurobiological solutions to the problems of sensorimotor control stand out:

1. Control is not assigned to a control center; brains do not have a central executive or planner or dictator. Control is widely distributed in the nervous system, though the secret of how this works has not yet been discovered.

2. In managing control, brains use *both positive and negative* feedback. This is uncommon in engineered control systems, partly because the combination often causes instabilities. The nervous system, however, combines them in a highly successful way. It uses positive feedback signals to predict what happens next in a feedforward control pathway, and negative feedback signals to make small corrections within the movement. This gives the system speed of response with minimum corrective wiggle.

3. Movement of sensory input systems, such as the eyes, appears to make certain computational problems of visual percep-

tion simpler rather than more difficult. For example, head and body movement help in determining depth of objects in a scene by creating motion parallax (near objects have greater relative motion than far, and objects in front of the fixation point move opposite to head direction, those beyond fixation point move isodirectionally). Hence animals often head-bob in order to extract more information through differences in relative motion of objects.

Body Movement Allows Representational Economies

In a surprising way, eye movement reduces the brain's computational load by reducing how much has to be represented in detail at any given saccade (small eye movements we scarcely notice but normally make about every 300 msec). As Dana Ballard, a computer scientist at the University of Rochester, points out, eye movement behavior takes advantage of the general stability of the world to economize on processing.[5] Ballard's idea is that only a small part of the visual field, roughly the central 2% (the foveal region), is processed to a high level.

On each saccade the brain samples a new 2% sector of the visual field. Attentional and motivational mechanisms appear to help guide the scan paths of the eyes. Psychophysical research shows that the eyes systematically scan a scene, returning many times to areas of high interest, relative to animal's current task. Ballard's hypothesis is that the brain does not need to have a richly articulated model of the whole world, because the world is out there to be sampled again and again. The world is largely stable, and relevant changes in the world scene can be picked up quickly.

Nonsmeary Vision During Head Movement

In neurobiology, an especially well-studied and revealing example of active sensorimotor control is the oculomotor system (figure 9). This circuitry is responsible for keeping the visual perception stable and crisp when head or eye (or both) are moving. But for its remarkable speed and tremendous accuracy, our visual image would smear every time we moved our heads. For a basketball player or a cheetah chasing a gazelle, this would be catastrophic.

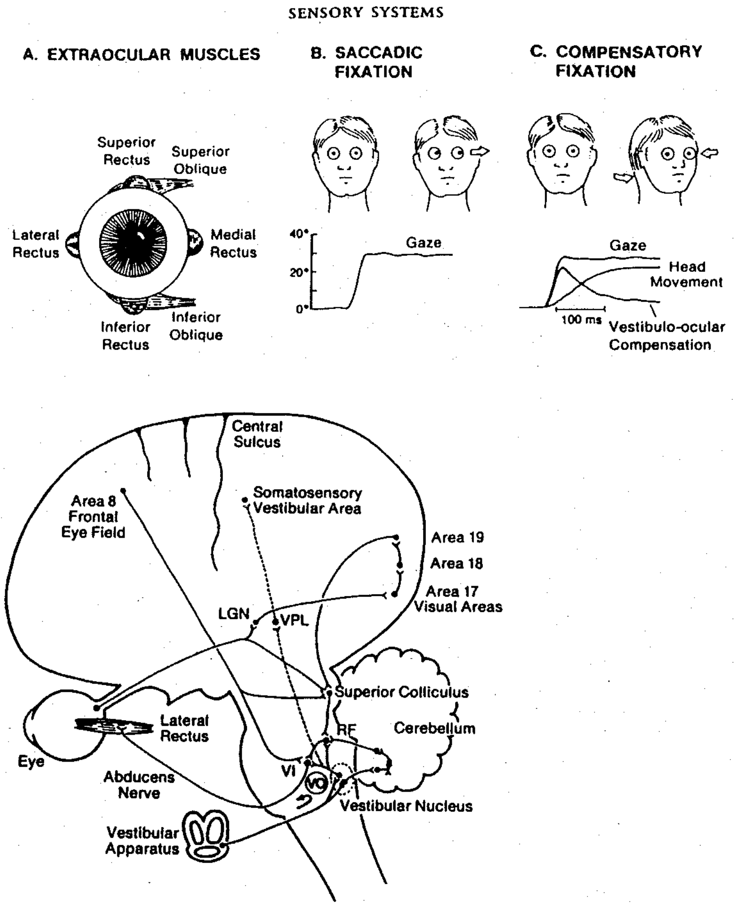

Figure 9. Pathways for visual tracking.[6] (Top) When the head moves, the perception of the object remains crisp as the eyes make compensating movements in the opposite direction. (Bottom) The input comes from the head acceleration-detectors in the semicircular canals. Information is processed in the vestibular nuclei, and the eye muscles are precisely signaled to contract so that the eyeball moves appropriately. (From G. Shepherd, *Neurobiology,* 2nd edition. Oxford: Oxford University Press, 1988)

Enough experimental detail about the anatomy and physiology of the oculomotor system is now known to support the interplay between dry modeling and wet experiments. This research is beginning to reveal how the oculomotor system can be so fast, how it modifies itself to accommodate structural changes in the eyeball, and how it makes efficient use of both feedforward and feedback signals.

Chips for handing the processing of head and eyeball movement are currently in the works at Caltech. Assuming continued progress in neurobiology and chip technology, a Mark I silicon "eye" that tracks slowly moving objects should be ready in a few years. An "oculobot" is still a long way from our fictional, full-fledged agribots, to be sure. Nevertheless, some of the ideas and technology leading in that direction are now in hand.

NEURO-REVOLUTIONS

We are on the brink of two "neuro-revolutions"—one in the science of the brain, and the other in the technology of brain-style computing. Growth of knowledge follows an exponential curve. Often, the more you have, the more you get—and the faster you get it. So it is with knowledge in neuroscience. Almost daily, surprising discoveries about the organization and mechanisms of nervous systems are reported. Setting neurobiological facts in a computational framework raises new questions about how networks of neurons work, which in turn give rise to productive, testable theories about how brains work—about how brains see, learn, and make decisions. In computer science, the VLSI revolution has provided us with unprecedented computational tools to transform what we know about the brain into silicon. Silicon retinae are in production, silicon cochlea are nearing production, and "oculobots" are on the drawing board. Although it is next to impossible to predict precisely other brainwaves in the genre, it is easy to forecast that ever more sophisticated neuro-engineering is in our future.*

*Portions of this article were based on our paper "Silicon Brains," Byte. October 1992

ENDNOTES

[1] Patricia S. Churchland and Terrence J. Sejnowski, *The Computational Brain* (Cambridge.: MIT Press, 1992).

[2] See I. B. Levitan and L. K. Kaczmarek, *The Neuron* (Oxford: Oxford University Press, 1992).

[3] Carver Mead, *Analog VLSI and Neural Systems* (Reading, Mass.: Addison Wesley, 1989).

[4] John Dowling, *Neurons and Networks: An Introduction to Neuroscience* (Cambridge: Harvard University Press, 1992).

[5] Dana H. Ballard, "Animate Vision," *Artificial Intelligence Journal* (48, 1991): 57–86.

[6] See also Stephen G. Lisberger and Terrence J. Sejnowski, "A novel mechanism of motor learning in a recurrent network model based on the vestibulo-ocular reflex." *Nature* (in press).