CHAPTER 1

BRAIN BEHAVIOR POINTS THE WAY

COPYRIGHTED MATERIAL

INTRODUCTION

The mysterious brain, underneath its understated facade, is just another computer, only with unimaginable parallelism and efficiency. It is a memory-based associative computer that works not with addresses, such as street numbers, but with features such as shape and shade. The end result is a sequence of mental pictures to light our way through life.

This book takes a fresh new look at artificial intelligence from the perspective of modeling human memory as a system, an all-digital system based on current knowledge about memory and cognition. At the bottom, neurons and their membranes are modeled with analog circuits to create arbitrary Boolean logic. At the top, memory system operations, including memorization, memory search, cognition, and learning, are modeled using digital circuits. The process of modeling is grounded in basic physics as expressed using standard analog and digital circuits. Standard analog and digital circuits offer a special advantage over the more abstract forms of modeling. Circuits may be integrated into silicon or simulated in software for systems that require some degree of artificial intelligence. Equally significant, brain models are a wellspring for new ideas and inventions.

For example, presented later in this book is a novel computer taken from brain memory models. As in neural circuits, signals switch in the kilohertz

Human Memory Modeled with Standard Analog and Digital Circuits: Inspiration for Man-made Computers, By John Robert Burger
Copyright © 2009 John Wiley & Sons, Inc.

range, which is relatively slow, but logic is massively parallel, and therefore the end result is actually more efficient than a modern dissipative machine. Once started, the logic requires practically no power, a feature suggested by the neural models in this book, and hence this novel design may rightly be claimed to be an adiabatic computer. An adiabatic computer is one whose heat dissipation approaches zero while executing programs, permitting a package smaller than would otherwise be possible.

It is hoped that this book will inspire applications, not just in computer design, but in related fields such as medicine, logic devices, robotics, artificial intelligence, neuroscience, and last but not least, education. Modeling is important to any science, but it is especially important to neuroscience, because these areas have a vast and exponentially growing quantity of data, the taking of which has immeasurably outpaced the development of models to explain and digest this explosion of information. Without adequate models, scientific progress would be difficult to impossible.

MODELING

It is amazing to realize that the performance of integrated circuits continues to increase exponentially, seemingly without bound. For example, the capacity of memory chips continues to double approximately every two years, thanks to the focused work of thousands of dedicated but nameless engineers and technicians. These people are involved in technology, which is really quite different from science. The two should never be confused. *Science* creates models in an attempt to understand nature; *technology* creates practical inventions to make life easier.

Simple accurate models are an indication of scientific maturity. For example, physics has a quantum model to explain observed behavior of particles. Unfortunately, this model is more mathematical than physical, but it accurately explains physical data. In the words of Werner Heisenberg (1901–1976), a founder of quantum mechanics, "the laws of nature which we formulate mathematically in quantum theory deal no longer with the particles themselves but with our knowledge of elementary particles."

Electrical science has James Clerk Maxwell (1831–1879), whose equations predict the behavior of electrical signals conducted through wires and radiated through space. It is perhaps obvious that electromagnetic waves are incorporeal, although they interact with matter in familiar ways. In modern physics, when describing waves and particles it may be concluded that for the most part, only the model is tangible; what is being modeled is not. Waves and particles, of course, make up everything that is possible in the physical world.

In the words of Herbert Alexander Simon (1916–2001), computer science is a science of the "artificial," that is, machines designed by engineers. Engineered computers are much simpler than brains evolved in nature. Modeling as pursued in this book is an attempt to create an emergent model of an evolved biological system that is concealed, well protected, and pretty much impenetrable. As long as human knowledge is finite no model in any field should be taken as a final measure of absolute truth. An important goal in working with all scientific models is that they may be changed and upgraded.

Modeling Goals of the Past

Models with standard analog and digital circuits have an advantage: They translate readily into engineering practice. In contrast, fashionable brain models usually do not have engineering practice as a goal. Consider three historical goals:

1. To study the sensitivities of a particular neural network to its parameters. For example, the output of a small rhythmic circuit in response to the timing of synaptic signals: Understanding can be increased this way, which in itself is considered worthwhile.
2. To discover some yet unknown relationship between a neural structure and its function. For example, why does the cerebral cortex have exactly six layers of neural gray matter? Modeling and simulation might reveal an answer to this question.
3. To discover which direction experimental investigation ought to take. Quite often, neural data are incomplete or nonexistent. For example, when a person makes a decision, there is some evidence that the brain makes the decision before the person is aware of the decision. Could it be that the brain tells you what to do, and not the other way around? This might be an interesting area of research.

Molecular biologists have strived to model neural systems molecule by molecule in what might be termed *bottom-up modeling*. Unfortunately, molecular models are very far removed from a brain system. Results can be pointless when building from the bottom up, with no master plan. Complex systems, both biological and artificial, necessitate not only bottom-up but also top-down modeling. *Top-down modeling* strives to synthesize known behaviors into a model, one that supports, or at least does not contradict, what we think we know.

Like computer systems, human memory systems necessitate both bottom-up and top-down modeling; standard analog and digital elements are an excellent choice for human memory. In this type of modeling it is useful to employ analog circuits to model the behavior of neurons, and digital circuits to model memory systems and cognition. Higher-level blocks in this text, in contrast to alternative models, are intended ideally to represent a particular interconnection of neurons, not just a floating abstraction defined by tables and calculus.

Brain modeling often involves higher levels of abstraction as supported by such simulators as the Genesis neural simulator and by the publicly available Neural Simulation Language. Unfortunately, as higher functions are modeled, the block properties of the components used in the simulation can depart from reality. For example, a block labeled "amygdala" might have no realistic structure at all, having instead a set of equations with no physical basis. On the other hand, if a "realistic" model based on individual molecules and ion channels is employed, computational complexity increases beyond reason, if not today's technology. Modeling like this is opposite the positive philosophy pursued herein: to illuminate a topic, not obscure it.

Standard circuit elements were never intended for modeling the paths of individual ions or electrons. An individual neuron can be modeled quite satisfactorily by assigning collective properties to electrical flow: for example, average current per unit time through a cross section of area. Whenever possible in this book, electrical models are derived from physical models. A simplified physical model, for example, may involve the average properties of thermally excited particles at temperatures known to exist biologically.

All models are, of course, based on a given body of knowledge. This implies that some information is excluded either deliberately or unintentionally. For example, the stochastic properties of ion channels are excluded as not being particularly useful in the model of a system. Modeling in this book is based primarily on the knowledge expressed in the appendixes and in the publications recommended for further study.

Uses of Models

Models involving circuits can be employed to reach historical goals, but models do more than this: They affect thinking. For example, it has been proposed that explicit long-term memory within the brain depends on synaptic growth. The problem is that growth takes time, whereas explicit long-term memories form rapidly, practically instantly, as can be demonstrated by those gifted with photographic memory. This fact brings the synaptic growth model into question, because it does not agree with what is seen in special cases.

In this book we present a model of long-term memory that does not require biological growth. Under this model, impressions may be captured by neural latches. Neural latches are not an enhanced form of short-term memory that eventually "time" out, but are completely different, latching instantly and holding indefinitely.

As another instance of models affecting thinking, it has been proposed that an action potential for cognition is coded not as a single bit, but with additional information. But a single bit from a single axon is sufficient for memory system logic, as demonstrated herein. So whatever burst duration, frequency, and amplitude may mean, they are not really needed for brain logic.

Natural philosophers, including Socrates (470–399 B.C.) have suggested that nature acts by the most economical means. Maupertuis (1698–1759) garnered attention with his *principle of least action* by observing that "Nature is thrifty in all its actions." This principle is apt for evolved systems, including neurons and brains. Under this principle, from the point of view of efficiency some models are better than others.

Relating to the synaptic growth model is the *one-neuron model*, in which a single large neuron is proposed to hold a given long-term memory. The problem is that no efficient circuit exists to send the contents of a single neuron into conscious short-term memory, which apparently is an integration of a great many features for a single image and a great many neural connections. This is an example of how the lack of an efficient circuit raises questions.

As another instance of the lack of an efficient circuit, consider a word of long-term memory. A word in this context is not a component of a language, but a collection of memory cells linked together by interneurons so that a memory can be recalled as a single image. It has been proposed that memories are added contiguously, implying that memory words keep expanding to some randomly long length. But the resulting model leads to an inefficient circuit with technical problems for searching, recall, and memorization, not to mention excessive duplication of feature-detecting neurons as common features are used over and over. Far more efficient for circuit-modeling purposes are words of approximately the same length, each with parallel connections to a common set of features. Thus, the lack of an efficient circuit model brings into question the idea of ever-expanding, randomly long memory words.

To summarize, simple models are important to simulation and prediction of behavior, and to general education. Models affect thinking in at least three ways:

1. Credibility increases when there is a model that predicts known facts.
2. Theories with simple models tend to be favored over those with complex models, all else being the same.
3. Some models are better than others from the point of view of circuit efficiency.

WHY THINKING DISSIPATES SO FEW CALORIES

Why does thinking use so few calories? Standard data suggest that a brain consumes roughly 10% of the net energy expenditure of the body and that it does so essentially without regard to level of mental effort. As evidence that brain neurons take precious little energy compared to muscles, consider the kilocalories dissipated for various activities, as shown in Figure 1-1. Note that these are "large" or "food" calories. One kilocalorie equals 4.184 kj of energy.

Averages in the data were taken over gender, age, ethnicity, and weight. The category of "reading" would include calories for eye movement and page turning, so the values above are for the entire body, not just for the brain and not just for thinking. Clearly, ordinary mental activities such as reading are nearly as efficient as sleeping. The brain rests very little, if any, even during

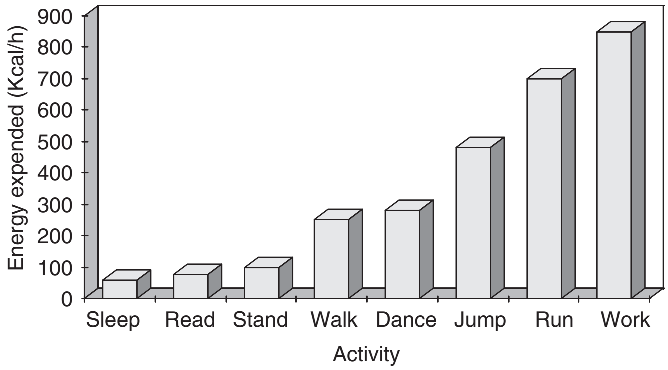

FIGURE 1-1 Typical calories expended. (Data from http://www.tooelehealth.org/Community_Health/CVD/Calories_Burned.html.)

sleep, as evidenced in part by dreams. Apparently, brain neurons, like those for breathing and heartbeats, need no rest.

An hourly use of 50 kcal/h translates into about 60 W of power. Applying the 10% rule, that leaves 6 W, more or less, for the entire brain. All joking aside, maintenance and all other brain activity absorbs less power than is expended be a dim light. If overhead energy is discounted, subtracting the calories used for maintenance, growth, and health in every cell of the body, few remain for the logical operations of neurons. There are nearly a trillion neurons in which virtually no energy is needed per neuron for ordinary mental activities, too few calories to measure easily.

It is remarkable that thinking uses so few calories, given that billions and billions of neurons are involved. Relatively simple man-made computers with only a miniscule fraction of the brain's computing elements, perhaps only a million gates, dissipate hundreds of watts and run very hot indeed. To convince yourself, simply touch a working computer chip. As a result, computers need heat sinks and cooling fans, and chips must be slow enough to minimize heat generation. Temperature poses a serious engineering limitation for computers.

Functional magnetic resonance imaging (fMRI) is an example of a modern tool (described in Appendix C) that appears to produce an image of the energy, or calories, consumed by different parts of the brain for given human actions. fMRI aims to observe oxygenated hemoglobin, related to calories consumed by neurons, but requires averages taken over several seconds to image signals that are extremely weak. The fact that fMRI signals are very weak is indirect evidence that calories expended in thinking are few.

fMRI probably does not image the extremely small energies dissipated for action potentials but may show a secondary effect related directly to action potentials. During an action potential one expects a brief reduction in heating within the neural membrane, the heating that occurs constantly for maintenance, growth, and health in any human cell. There normally is about

-70 mV across the membrane, which has a small conductance and so dissipates a small amount of energy. During an action potential the *average* voltage across a neural membrane is slightly lower because membrane voltage alternates between about -70 and $+40$ mV. This reduces dissipation in the membrane and probably changes the amount of oxygenated hemoglobin being processed, as fMRI purports to observe.

Common sense tells us that mental exercise is not a way to lose weight. Mental activity does not cause one to work up a sweat, nor does it get the heart racing. It appears that no energy at all is required for ordinary mental activity, including such items as recall, memorization, reasoning, reading, and listening. Mental activity is way down on the kilocalorie scale, not counting energy for cellular maintenance, growth, and health. Learning in the memory system model falls into a different category since energy is expected to be required for synaptic growth.

The lack of energy dissipation for ordinary mental activity is more than an interesting topic for dinner table conversation, because it points to the adiabatic model of neurons as logic devices, those that operate with little or no dissipation of energy. Aside from the adiabatic neuron, other adiabatic models for computation may be found in the theory of quantum mechanics.

Computer engineers may be interested to know that adiabatic logic is theoretically possible using common solid-state circuits. Once charged with electrical energy, such circuits transfer charge back and forth without loss, eventually returning all charge and its associated energy back to the power supplies from whence it came. With careful design, adiabatic logic leads to adiabatic computers for packaging into very small volumes. Small volumes are feasible because adiabatic logic runs cool. If the brain is any inspiration, we need to greatly expand our use of adiabatic logic for everyday life.

THE MIRACLE OF PARALLEL PROCESSING

The human brain is unique in nature in that it supports millions and millions of channels computing simultaneously. They are serviced by a large number of sensory inputs and deliver a large number of motor nerve outputs, all in parallel. The strength of a system like this is easily overlooked, since it operates effortlessly. How does one comprehend millions of little computers in parallel? We humans are lucky to juggle two or three things at once.

Mental activity is generally sequential: that is, one thing at a time. But memory searches are not sequential. Effectively, a human being can poll all memories in parallel and recall a particular episode from millions of events experienced long ago. This is a type of parallel processing. A style of man-made parallel processing known as *associative memory* works along these lines but on a much smaller scale. Associative memory finds items not by addresses, but by attributes such as patterns in the data. Modern computers employ associative memory only to an extremely limited extent: for example, to quickly find the most recent addresses

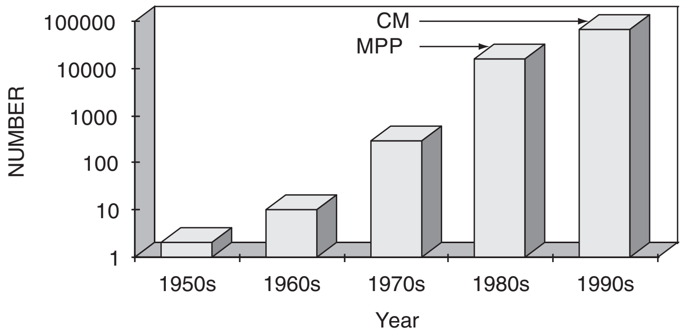

FIGURE 1-2 Brief history of massively parallel (associative) processing (CM = connection machine; MPP = massively parallel processor).

used by the microprocessor. To emulate brains for use in robots and artificial intelligence, we need to greatly expand our use of associative memory.

The possibilities for problem solving increase dramatically if millions of basic computations can be carried out simultaneously. Unfortunately, many people fail to see the possibilities of massive parallelism, because it requires a different type of thinking. As an example of a new way of thinking, consider finding prime factors. If a given large number is copied many times, and concurrently, if each copy is divided by a different prime number, the prime factors of that large number can be identified instantly. Simply flag those results whose remainder is zero.

Autistic savants perform amazing mental feats, including the immediate recognition of a prime number (one that cannot be divided evenly except by itself and unity). We do not know how they do it, exactly. One theory is that they simply memorize. However, given all that a savant can do, simple memorization is unlikely. Parallel processing must be occurring.

It may be noted that parallel processing is not limited to arithmetic. Finding the shortest path through a given maze, for example, is possible using a massively parallel processor, but only if the maze is not too large. Although not yet involved in everyday applications, the idea of massively parallel processing is certainly not new. Figure 1-2 is a snapshot of historical efforts at parallel processing.

The size of parallel processors tends to be increasing, yet parallel processing remains a topic of research and has run into serious problems. People generally prefer sequential procedures, which brings us back to the statement that parallel programming requires thinking in unaccustomed ways.

SINGULARITY

Machine intelligence approaching the human level is expected to produce a step in the economic growth rate of the past, as nicely exposed in an article by Robin Hanson (1959–). According to Hanson (*IEEE Spectrum*, June 2008, pp. 44–50)

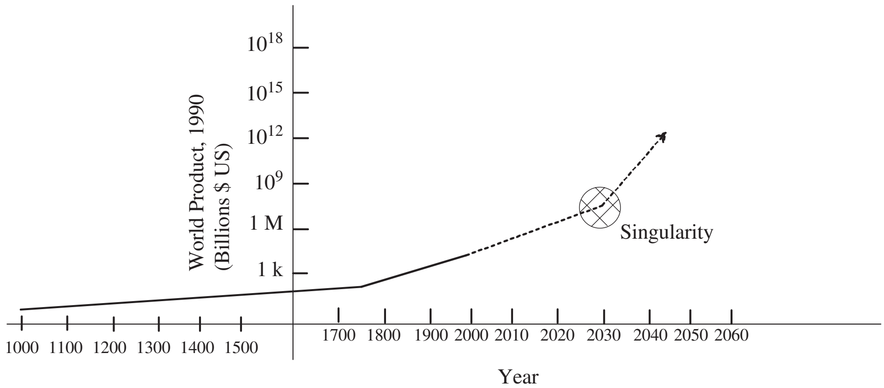

FIGURE 1-3 Idealized economic growth, including a singularity.

and others, a societal discontinuity is imminent, brought about by robots with artificial intelligence, similar to the original industrial revolution. The date of this new robotic revolution is uncertain, but it is known as the *singularity*. This word comes from an idealization of economic output as shown in Figure 1-3. Mathematically, the rate of change of economic growth, the slope of the curve, will experience an abrupt step. Going even further mathematically, the rate of change of the ‘‘rate of change’’ will be singular, that is, go to infinity at this magical point.

The original industrial revolution has been going on for quite some time, since about 1750, a general result of the pervasive application of machines. Since then the world economy has doubled roughly every 15 years or so. The switch from an agricultural society to an industrial society resulted in a speedup of roughly 100. Assuming a similar speedup as a result of the singularity, the economy would double not in 15 years, but in a month or two, assuming no setbacks due to war, plague, or cosmic disaster.

We see the original industrial revolution, but do we understand it? Machines are partially responsible. Machines like the steam engine, were possible because of advances across a broad front in metallurgy, metalworking, and engineering knowledge. No single spectacular gain in one area is responsible, since a complex machine has thousands of parts, each of which might depend on dozens of technologies. A gain in one technology causes only a slight improvement overall. This is the *law of diminishing returns*. To have an economic revolution there must be exponential growth.

Exponential growth was provided in the eighteenth century by capitalism, in which incremental improvements and novel applications were accomplished by countless hands-on builders. The original industrial revolution continues today sustained with higher-capacity memory chips, improved computers, and the World Wide Web, as promoted by thousands of innovative capitalists.

What is there today sufficiently broad-based as to induce a singularity? We know that improvement in just one sector of technology is insufficient, owing to the law of diminishing returns. But what if there were something that corrected a chronic shortage in all sectors: human attention and intelligence. Most financial gain in rich countries today goes for direct and indirect costs of labor. An innovation that drastically reduces this cost could very well start an economic revolution.

Machine were involved in the original industrial revolution; intelligent machines may be involved in the second, this being a cornerstone of the singularity hypothesis. But how can a machine be intelligent? One answer is that machine intelligence will follow shortly after computer hardware approaches the performance of a human brain. In an attempt to approach what the brain does, one approach is to ‘‘reverse engineer’’ the brain with the aid of scans and modeling. When such processes are perfected, a particular brain might be manufactured. The duplicated brain, of course, can always be unplugged, unless it is adiabatic and requires no external power.

Modeling efforts are currently under way. Project BlueBrain is a joint effort by IBM (International Business Machines) and École Polytechnique Fédérale de Lausanne in Switzerland to reverse-engineer a brain, but not a human brain. Their goal is mapping and modeling the roughly 10,000 neurons and 30 million synapses in a rat’s neocortical column. The real goal, by the way, is a human brain. If the mental powers of a human could be approached, it would be more than a scientific curiosity. The chief factor of economic production that has been chronically scarce throughout history, human intelligence, would suddenly become widely available with mass production.

Imagine that machines approach human cognition and are able to perform most human jobs. With a growing workforce of intelligent computers created more quickly than it takes to breed, raise, and educate humans, the economy would explode. The cost of producing such a workforce will drop at an accelerating rate. Intelligent computers may even learn to design and manufacture other computers, all of which suggests a step in economic growth rate.

What does the future hold? To quote Robin Hanson: ‘‘Stuffed into skyscrapers by the billion, brainy bugbots will be the knowledge workers of the future.’’ Operating at machine subsistence levels, their main cost will be rent for their miniature volumes, occasional parts for repair, and hopefully, no air-conditioning bills. The singularity is a plausible view of the not-too-distant future—and it is not all bad. Humans would labor less for income while gaining by renting real estate for intelligent machines and by investing in the maintenance of computers. They would have ample time to labor for pressing matters such as exploring the universe.

THE BENEFITS OF READING THIS BOOK

This book is written primarily for those with interests in artificial intelligence, computer science, neuroscience, robotics, and peripheral fields such as artificial neural networks, psychology, and medicine. The book views the brain from a circuits and systems perspective, circuits and systems being the author's major at the University of California–Los Angeles in the 1970s. Knowledge in the aforementioned fields is enhanced by models of human memory using standard analog and digital circuits. Beyond formal knowledge, there is much in this book to help a reader personally. For example, when trying to spark the retrieval of something that has been forgotten, it is helpful to think of a wide variety of cues. When attempting to recall a person's name, we use such clues as the name of his or her spouse or the type of car the person drives. Sometimes, for example, a forgotten name pops forth immediately upon seeing the person from whom you first heard the name.

Nearly everyone has tried to remember something, but cannot; hours later the desired memory emerges, usually at an unexpected moment. This is because long-term memories are routinely searched in the background, subliminally. Humans are not particularly aware of this phenomenon of *delayed recall*, but once they are, they may take advantage of an everyday process in their own brains, the process of subliminal memory search. To recall facts, make decisions, or solve problems, it is wise to allow the unconscious to work for you, giving it time to do so for a day or two.

Dreams, daydreaming, or brainstorming are common experiences. In dreaming, a person becomes aware of random episodes from long-term memory, quite often relating to the solution of a difficult problem. The brain never quits, it seems, trying to solve a difficult problem, perhaps a problem with emotional impact that has no logical solution. Such common experiences as dreaming and delayed recall show that memory searches proceed randomly, unnoticed. When a sought-after memory is found that matches all available cues, it pops immediately into your conscious short-term memory with a certain modest excitmemt.

Once it is decided that something will be placed into unconscious long-term memory, it also helps to memorize cues for retrieving that information. Also, one must rehearse in a regular way both information and cues for memory; rehearsing in a random or irregular way is not expected to be as effective in the model of this book. Regular rehearsal is necessary to trigger a memorization enable signal, after which information in short-term memory goes automatically into the next available location in subconscious long-term memory.

Learning is facilitated by understanding the learning models discussed in this book. Placing information from the senses or from long-term memory into short-term memory on a regular basis is detected internally as a "need to learn." Regular practice is important at first because this enables the brain to generate a need to learn.

Upon receipt of a need-to-learn signal, there is a decoding of the contents of short-term memory to spur the growth of neurons and synaptic connections where they are needed. Over a period of time, this growth fosters new abilities, such as learning to recognize a new feature without analysis, or learning to follow a complex procedure automatically in a mindless way. At first, new synapses are small and could dissolve in a couple of weeks because of thermal and chemical activity. Therefore, it is helpful for a person to engage in overlearning and relearning to resist forgetting. Relearning is usually easier than the initial learning; it is encouraging to realize this progress.

A person can "learn how to learn." This is largely a matter of organizing the material to be learned into a more efficient form. For examples, numbers to be learned can be imagined as sporting event scores; audio tones and visual textures can be given names; a message to be learned can be broken into phrases that rhyme; the steps of a procedure to be learned can be numbered to facilitate learning.

Of course, one must always keep in mind that learning involves synaptic growth and is totally different from memorizing, which involves a latching of neural memory cells. Memorizing is intended for recall through the apparatus of short-term memory. Learning can essentially bypass short-term memory so that what is learned can be realized automatically without thinking.

Rehearsal aids learning; if a person wants to learn a skill, he or she must practice. Synaptic development requires time; so if nothing else, a person must patiently practice what was learned and to allocate sufficient time to reinforce what was learned.

Last but not least, neural circuit models might provide a professional benefit to medical practitioners. By understanding neural circuits, a brain malfunction might be related to circuit faults, which gives a practitioner the language to describe the probable cause of a malfunction. Subsequent circuit simulation might suggest corrective actions and might someday even predict the outcome of proposed surgery prior to an operation. Better medicine and better engineering are examples of real physical benefits that many are pursuing.

OVERVIEW OF THE BOOK

Standard circuit models are provided throughout to avoid the usual vagueness when it comes to describing brain operations verbally. Maupertuis gave us a principle of least action, interpreted to mean that systems evolve to be efficient. This principle has been applied to the modeling of neurons as logic devices all the way up to system models for associative memory. This principle leads us to a model of the brain as an adiabatic, massively parallel computer that does not lose information. Just the thought of an adiabatic, massively parallel computer is enough to suggest revolutionary new ideas in the minds of inventive engineers. Next we describe how the book evolves.

Chapter 1: Brain Behavior Points the Way In this chapter we introduce neurons in parallel that require virtually no energy for signals. Their signals are modeled to exist above the calories dissipated for common maintenance, growth, and health, common to all cells in a mammalian body. This adiabatic model creates new possibilities for explaining brain behavior. For example, an adiabatic model like this, in which action potentials come and go without dissipation, creates the possibility of long-term memory based on a simple circuit.

Chapter 2: Neural Membranes and Animal Electricity A voltage differential is developed because of charge transfer through a thin neural membrane, and this differential places it under significant electrical stress. To model a neural pulse electrically, sensitive regions of a membrane are modeled as ferroelectric: that is, as sensitive to an electric field. Reducing the electric field triggers a charge transfer that is driven by thermal energy. The resulting pulse is regulated by the ferroelectric particles within the membrane in that as internal voltage accumulates, a reversed electric field forces a reversal of the sensitive particles within the membrane. This reversal initiates a decrease in internal voltage. As voltage drops below its original equilibrium value, the pulse is forced to terminate because ferroelectric particles are driven back into their original positions. The sensitive particles of interest within the membrane are reset, and this enables a return to rest conditions. The model is such that an unlimited number of additional pulses can be triggered in a continuous manner if necessary.

Chapter 3: Neural Pulses and Neural Memory In this chapter we employ a simple physical model involving thermally active ions to derive by hand the waveform of a neural pulse. Underlying this model is the probability that thermally excited ions and stray electrons can tunnel into the sensitive regions of a membrane. Tunneling explains the magnitudes of the charge transfers commonly observed.

To model short-term memory neurons, ionic concentrations are modified to create and then hold positive charge within dendrites that have been exposed to excitatory neurotransmitter ions. This causes continuous triggering of the soma and axon with enhanced frequency but reduced amplitude lasting for a few hundreds of milliseconds. Connecting neurons respond to this signal. Explicit long term memory, in contrast, is modeled as a latching that occurs when the output neurotransmitters of an adiabatic neuron are fed back to dendritic receivers. Once latched, memory cells hold their features indefinitely without significant energy dissipation.

Chapter 4: Circuits and Systems for Memorization and Recall An all-digital model of the memory system is synthesized with an eye toward explaining the brain technically. Central to this model, which is a rudimentary cognitive architecture, is short-term memory organized as a long word of features. Short-term memory, that of which a person is conscious, is connected to millions of similar words forming a large associative memory which is assumed to hold all relevant information, including past problem solutions and decisions. This memory is searched by random selection from short-term memory, cues that are applied to recall long-term memory words up to tens per second. Such recalls, alternated with sensory images, are flashed subliminally to an interface circuit associated with short-term memory. Here, an encoder, such as a priority encoder, calculates an *index of importance* as an ongoing process. If the index for an image is higher than the current index for short-term memory, that image is gated immediately into short-term memory. This establishes a *direction of attention*, essentially a moving picture in conscious short-term memory as supported by a model based on standard digital logic.

Chapter 5: Dendritic Processing and Human Learning Dendritic pulses created by excitatory neurotransmitter ions are recognized to be electrical solitons that propagate away from their point of creation without dispersion or attenuation. Simulations indicate that solitons are easily reflected at the soma, and when they collide with oncoming solitons, they annihilate each other, thus reducing the number of solitons triggering the soma. It is shown that solitons are capable of arbitrary Boolean logic. AND–OR logic is possible, as solitons charge soma capacitance to trigger a neural pulse. Neurons with digital properties are shown to support an all-digital learning model for humans. Two types of learning are identified: (1) combinational learning, to recognize a new feature in sensory data without having to stop and think about it, and (2) state machine learning, for new procedures that are executed without concentration, such as dancing or reciting a poem. Circuit models are suggested in support of all-digital learning for both new features and new procedures. Both types of learning avoid passing information through short-term memory for evaluation, thus increasing efficiency as necessary for species survival. Standard digital logic is modeled in support of the ideal of an intelligent robot that can actually learn beyond mere memorization.

Chapter 6: Artificial Learning in Artificial Neural Networks In this chapter we present learning as defined for artificial neural networks. Artificial neural networks use analog weighting factors and analog summation, so differ from the all-digital ideal. Learning in an artificial neural network is equivalent to designing weighting factors, often a tedious iterative process. Overall, the tremendous success of artificial neural networks demonstrates the importance of learning in machines.

Chapter 7: The Asset of Reversibility in Humans and Machines After considering the abilities of gifted savants, we speculate that their advanced abilities are the result of parallel processing within long-term memory, a type of processing in which no energy is dissipated. It is shown that no energy is lost in a system that is *electrically reversible*, that is, in which a given amount of charge is applied, all of which is recovered. To be electrically reversible, a computer must lose no information and so may be logically reversible, although a logically reversible machine need not be adiabatic. A reversible programming concept is introduced based on a “wiring” diagram as envisioned for the savant brain.

Chapter 8: Electrically Reversible Nanoprocessors In this chapter we present a case study for the original design of an adiabatic parallel computer using solid-state technology. Brain-inspired, the design avoids data buses during a computation—a major bottleneck and a source of heat dissipation in conventional computers. Central to the design are nanoprocessors, words of memory with conditional toggling capability within each word. An array of nanoprocessors constitutes an associative processor that can be designed to be electrically and logically reversible. Example programs in a wiring diagram are provided for vector addition and subtraction.

Chapter 9: Multiplications, Divisions, and Hamiltonian Circuits In this chapter we demonstrate a variety of programs possible for an electrically reversible parallel computer, including multiplication and division based on add-shift and subtract-shift algorithms. Solutions to SAT or NP-complete problems using a reversible parallel computer are introduced. The problem of finding all Hamiltonian circuits in a small graph is discussed, particularly the easy part: checking to determine if a given cycle is indeed Hamiltonian. Electrically and logically reversible computers are limited in practice by the number of nanoprocessors that can be brought to bear on a problem. Molecular-sized nanobrains, if they become available, will significantly increase the number of nanoprocessors, although perhaps not enough to solve SAT problems with thousands of variables. Qubits can be smaller than nanobrains and show promise for solving large problems.

Chapter 10: Quantum Versus Classical Computing The goal of this final chapter is to introduce quantum computers and to identify what the various biological, electrical, and quantum computing systems have in common. All such systems may use wiring diagrams, for example, so the transforms implied by reversible gates such as UN, SCN, DCN, and MCN are held in common. Wiring diagrams are used to explain savant brains as well as to program adiabatic parallel computers and quantum algorithms and to manipulate qubits, these being physically reversible. Such systems have in common that no information is lost and that basic computations are adiabatic, not counting energy overhead to maintain a workable environment.

The following appendixes may prove useful for general information.

Appendix A: Human Brain Anatomy This is a summary of basic brain information, some of which is used by the book's memory model.

Appendix B: The Psychological Science of Memory This is a summary of basic memory psychology, some of which is used by the book's memory model.

Appendix C: Brain Scanning This is a summary of fundamental imaging and scanning methods used to study the brain.

Appendix D: Biographies of Persons of Scientific Interest This is a short collection of biographies of interesting and occasionally unpleasant characters who contributed to topics discussed in this book.

For Further Study Finally, there is a brief list of published material for further study.

APPLICATIONS OF THE MODELS IN THE BOOK

Unlike models with abstract symbols, those expressed as standard analog and digital circuits can be simulated with ordinary software, constructed with everyday technology, tested by standard methods, and applied in various ways. The models described in this book bring to mind applications that today's students will undoubtedly come across in the course of their careers.

Artificial Membranes

One goal in the field of artificial membranes is to build a membrane that is as close as possible to a biological membrane using nanotechnology and genetic engineering. Toward this end, the models in this book provide an interesting perspective. Here we model membranes as containing sensitive regions with ferroelectric properties such that membrane particles are held together tightly by an electric field, but relaxed in a lower field. When relaxed, or triggered, it is possible to have charge transfers that constitute a neural pulse.

An artificial membrane ideally would behave like this, although such a membrane has yet to be manufactured. Interestingly, purple membrane films have been found to display ferroelectric behavior (termed *bioferroelectricity* in the literature). Purple membranes, characteristically hexagonal in shape are two-dimensional structures consisting of a transmembrane protein surrounded by 10 lipid molecules. They are relatives of man-made liquid crystals (a 5 billion market) based on ferroelectric properties.

Practical membranes have important applications, such as microfiltration, reverse osmosis, pervaporation (separation of liquids by vaporization), gas separation, dialysis, and chromatography. This translates into water purification, removal of microorganisms in dairy products, water desalination, dehydrogenation of natural gas, hemodialysis, and fuel cell components.

For the most part, artificial membranes with ferroelectric properties are not yet available, although many proposals have been made. For example, one proposal is to engineer membranes in the form of cells with internal bipolar charges within the voids, to mimic ferroelectrics. Another interesting proposal is to build a silicon chip with patches of artificial membrane so that filtration might be regulated using standard complementary metal–oxide semiconductor (CMOS) electronics. The neural model in this book is one way to gauge artificial membranes. Such membranes may be characterized by the pulses they generate when immersed in ionic solutions and adjustable electric fields.

If an artificial membrane approached the behavior of a neural membrane, it would facilitate molecular-sized elements known as *nanodevices*. Currently, a variety of these have been proposed to assist with cancer detection, diagnosis, and treatment. What is needed are general-purpose nanodevices: not just as nanobots for medicine, but for general applications to computers and communications.

Imitation Neurons

Imitation neurons approximate biological neurons both physically and electrically and are unrelated to *artificial neurons* for artificial neural networks. Artificial neurons are merely electronic amplifiers or computer subroutines to simulate amplifiers. Imitation neurons are modeled after biological neurons and might someday be fabricated using the methods of *nanotechnology*: engineering on a molecular scale, normally 1 to 100 nm, and to the fabrication of devices within that size range.

As decades roll by, analog and digital amplifiers are shrinking in size and increasing in efficiency, slowly approaching the efficiency of the everyday neuron. Beginning with vacuum tubes early in the twentieth century, amplifying technology evolved to the transistor in midcentury. Today's transistors are much smaller and more efficient in terms of heat dissipation, with temperature increases well below the melting point of silicon thus permitting very large scale integrations that were impossible a short time ago.

We have now entered an age of nanotechnology that promises further decreases in device size. Nanotechnology has arrived because of the recent availability of such novel tools as the atomic force microscope and the scanning tunneling microscope. Combined with refined processes such as electron beam lithography and molecular beam epitaxy, the deliberate manipulation of nanostructures has become possible. The birth of nanotechnology is generally assumed to be in 1989 when IBM scientist Don Eigler wrote out the company's logo using 35 individual xenon atoms arranged on a nickel plate at low temperature and high vacuum.

Imitation neurons are a goal, and once perfected, would be a superior choice for nanoprocessors of the future. Imitation neurons are expected to be small and flexible, like biological neurons. Currently, flexible circuits are in great demand for devices such as connectors, liquid-crystal displays, and digital cameras. A minuscule neuron, in principle, can generate any large-scale Boolean function or any large-scale analog-to-digital conversion for an astonishingly large number of inputs. Most important, properly constructed imitation neurons would merely borrow energy from the ionic solutions in which they are immersed, like the biological models of this book. They would produce a burst of pulses internally, dissipating neither energy nor power.

As a sign of progress in nanotechnology, IBM scientists recently announced that they have created an embryonic nanoprocessor using a single carbon-nanotube molecule and a standard semiconductor processes. The circuit, called a *ring oscillator*, consists of 12 field-effect transistors laid along a carbon nanotube 18 μm long, which is about one-fifth the width of a human hair. Clearly, it is far from molecular sized, but remember, this is only the beginning. The direction being pursued is to make circuits faster and compatible with regular silicon technology for everyday integrated circuits.

At the molecular level, random thermal activity would be quite rough on rigid technology, so flexible wet technology must be considered. In the context of watery technology, it is conceivable that imitation neurons might someday replace damaged biological nerve connections. This is not a surprising concept in view of successes in brain–machine interfacing. Neuroscientists have significantly advanced brain–machine interface technology, to the point where severely handicapped people who cannot contract even one leg or arm muscle can now compose and send e-mails independently and operate a television set. They are using only their thoughts to execute these actions. One day these and other handicapped persons may be able to feed themselves with a robotic hand that moves according to their mental commands. The hope is that imitation neurons might enable the muscles of the paralyzed to be useful again.

Artificial Neural Networks

Artificial neural networks are based on artificial neurons, each of which is composed of a weighted sum and a comparator, giving a true or a false output. An artificial neural network can be taught to recognize important patterns in a large field of data using an iterative algorithm to design the weights. The fact that a machine can learn this way is quite amazing.

Applications of artificial neural networks include system identification and control (vehicle control, process control), game playing and decision making (backgammon, chess, racing), pattern recognition (radar systems, face identification, object recognition), sequence recognition (gesture, speech, handwritten text recognition), medical diagnosis (tumor recognition), financial applications (recognition of trends in the prices of stocks), data mining (knowledge discovery in databases), visualization, and e-mail spam filtering. Artificial neural networks are undoubtedly an engineering success.

Artificial neural networks employ analog weighting factors and linear summation and thus depart from the all-digital ideal. Digital circuits are ideal since they are small, inexpensive, and tolerate noise. The all-digital system model in this book assumes thousands of neurons uniquely detecting thousands of features, with each neuron essentially an arbitrary digital circuit with many inputs. Many billions more constitute an associative memory with processing ability, an associative processor.

Learning is exceedingly important to the illusion of intelligence. Learning in the book's model begins with a need-to-learn signal that permits the directed insertion of new digital circuits where they are needed, somewhat as in a field-programmable gate array.

Beyond artificial neural networks, pattern recognition can be done in other ways, assuming the availability of parallel nanoprocessors. For example, if analog information has been converted into a sea of digital data in no particular order, the data can be searched with ease. Words with a given pattern can be located instantly by parallel nanoprocessors efficiently solving "needle in a haystack" problems, provided that the haystack is not too large. Inspired by models of human memory processing, designs for electrically and logically reversible parallel nanoprocessors are given later in the book.

Computer Design

Norbert Wiener (1894–1964) defined the original meaning of the term *cybernetics* to be the study of control and communications in humans and machines. There seems to be no doubt that the study of control and communications in people and machines has been an inspiration over the years for various computer designs. The amazing brain has always been an inspiration for computer design, beginning with the fact that both brains and computers are memory based. Louis Couffignal (1902–1966), another pioneer of cybernetics, characterizes cybernetics as "the art of ensuring the efficacy of action." This characterization brings to mind the guiding principle of modeling with standard analog and digital circuits: Maupertuis's principle of least action.

Historically, brain models were not all that accurate, but they still served the field of computer design. Models help inventers in mysterious ways. For example, neurons and neural systems are modeled as electrically reversible, implying that no information or energy is lost as a result of cognitive activity. This interesting model applies readily to computer design. Once an engineer knows enough to conserve charge, computers are easily designed to be adiabatic with no power or heat dissipation, as presented in this book, using CMOS as an example technology. When little or no heat is dissipated, engineers gain an option to design computers into much smaller packages without concern for temperature increases resulting from heating.

To compensate for the slowness of adiabatic logic, massive parallelism is desirable, as in human memory. This implies billions of little nanoprocessors all operating at once, analogous to words of biological memory, synchronized by signals from peripheral registers, analogous to short-term memory. Data buses, essential to conventional random-access and read-only memory (RAM and ROM), are avoided, since they create bottlenecks and waste energy. Once developed, massively parallel adiabatic computers will find uses in areas of space exploration, medical implants, databases, human guides, and computations that are impossible in any other way. As molecular-sized nanoprocessors materialize, their numbers will increase beyond belief, helping to solve difficult problems that have thousands of variables.

Robotics

Although the appearance and capabilities of robots vary vastly, all robots relate to a movable mechanical structure under some form of autonomous control. A sophisticated robot has some degree of artificial intelligence, or ability to make choices based on the environment, often using a preprogrammed sequence. This conveys the sense that a robot has intent or agency of its own. Robots that are more or less intelligent are currently being used in manufacturing, lifting and moving in distribution centers, household chores, military operations, and perhaps most important, in toys.

Cognition models may be visualized for robots to enhance their autonomy. For instance, for good performance, sensory inputs and motor outputs should operate in parallel. The human memory model in this book uses a word of short-term memory in conjunction with a very large number of long-term memory words, equally wide. This structure facilitates associative processing based on images in short-term memory. Microprocessor serial processing, still in use for most practical robots, is poorly suited to a multidimensional environment.

Ideally, sensory information would go into short-term memory as it does in a human brain. As in the cognitive model presented in this book, cues for an ongoing memory search may be taken from short-term memory. Pseudorandom combinations of cues may be employed to recall memorized images to a subliminal level at rates of perhaps tens per second. Each recall may be analyzed for importance according to specified criteria, where importance is a digital encoding function. Sensory images are analyzed similarly. When an index of importance surpasses that of the current contents of short-term memory, which is fading away, a new short-term memory enters. This creates a moving picture or a direction of attention in short-term memory.

If a robot is to appear intelligent in a minimal way, it must be able to make helpful decisions. When a decision is needed in a system of the digital variety, a memory search recalls similar problems to locate a ready-made decision. If a problem cannot be solved logically in this way, attempted solutions, taken randomly, will make a robot seem all too human. A robot that identifies a problem and recalls a procedure to solve the problem certainly has a degree of artificial intelligence, but this is minimal intelligence. It is necessary to learn to behave better. The next step in the evolution of robots is to enable them to "learn" procedures without having to bring every step of a solution into short-term memory. Humans routinely perform procedures to solve problems without full evaluation in short-term memory; this is possible because of a process of learning in which *neural state-machine learning* develops within long-term memory. Neural state-machine learning permits us to walk, brush our teeth, or memorize a long poem without an excessive amount of pondering. The ideal robot could do this, too.

Humans also learn to recognize special combinations of sensory inputs without having to ponder a number of related combinations. For example, a human can learn to recognize a special color such as chartreuse, a mixture of yellow and green. Humans need not tie up their short-term memory cells with items relating to yellow and to green, and they do not necessarily need to recall a dictionary of color names. They can learn to recognize chartreuse *immediately*, termed *combinational learning* in this book. The day is coming when robots, in the interests of efficiency, will have the capability for combinational learning.

Using the all-digital model described in this book, a *need-to-learn signal* is first created by digital filters associated with short-term memory. Learning is accomplished with decoders that connect to where additional logic circuits are required. Unlike the field-programmable gate array, human learning does not need a download from an outside computer. Learning is self-contained. Suffice it to say that without learning, robots will always be awkward, and except for the simplest behaviors will black-out as their memories are tied up unnecessarily with mundane computations.

Artificial Intelligence

Artificial intelligence is a branch of computer science usually involving software for mainframe computers, to perceive important aspects of the environment and display a reasoned, humanlike response. Among the topics of interest in artificial intelligence are reasoning, knowledge, learning, and the ability to communicate. A high level of artificial intelligence has not yet been achieved and is considered a goal for the future.

Artificial intelligence has been approached as a programming problem. A traditional measure of artificial intelligence is to send questions to a computer via a keyboard and to judge by the answers whether or not there is a human at the other end. This is the *Turing test*. As an example of what has been considered artificial intelligence, Deep Blue was the first computing machine to win a chess match against a reigning world champion, Garry Kasparov. Other examples of artificial intelligence are the ability to understand sights, such as faces, and to understand sounds, such as spoken commands.

The field of artificial intelligence could benefit from the model of human memory and cognition presented in this book, especially artificial intelligence as it applies to self-contained robots. Brain circuits operate in parallel, as opposed to the high-speed serial processing that computer manufacturers prefer to sell us. Engineers are well paid to increase clock speeds in order to run long programs of serial instructions. Ironically, the result is a machine that is sometimes slower than older models, because software always grows faster than hardware. The point is that parallel processing will be required to match what humans do routinely.

Of all the aspects of artificial intelligence, the ability to learn is a most convincing sign of intelligence. One aspect of learning involves the development of neural state machines directly embedded in associative memory, with the capability to execute procedures without evaluation in short-term memory, as noted above for robots. It would be nice if a machine could learn to fetch a ball, like a dog does. Better yet, perhaps a machine could, without a lot of additional programming, learn to walk across a garden while pouring a glass of wine.

Intelligent machines need to recognize sights and sounds without saturating a central processor. For example, it would be nice if a machine could learn to recognize a new face that appears often, such as that of a household pet. The learning involved is called *combinational learning*, another form of all-digital learning explored in this book.

Neuroscience

Neuroscience is the scientific study of the structure, function, and development of the nervous system, so traditionally, it is a branch of biology. Neuroscience expanded significantly in the second half of the twentieth century and now includes branches in molecular biology, artificial neural networks, and computational neuroscience. Here we emphasize electrical models, so apparently "electrical neuroscience," although seldom mentioned, is another branch of neuroscience.

The neuron in this book is modeled as an analog circuit. These analog underpinnings suggest that in its own way, a neuron can calculate arbitrary Boolean functions. Using many such gates, a memory system can be modeled to provide neuroscientists with a new perspective. In this book, problem solving and decision making as well as information retrieval are considered to be memory related.

Modeling with digital circuits clarifies the actions taken by a memory system when cues are inadequate or when cues are ambiguous. Modeling offers a novel point of view for most traditional issues in memory theory, such as strong memories versus weak memories, accuracy of memory, and threshold-of-effort for memorization and the meaning of recognition. All of these central issues are touched upon by the book's human memory system model based on standard analog and digital circuits.

The Beginnings of a Cognitive Architecture

Random searches are useful in computers, so here is an idea from computer science that might be useful in practice. As modeled in this book, but now with a little more detail, to recall something, a name, for example, cues in short-term memory are made available. But when a name cannot be remembered, a memory search proceeds in the background, subliminally, and may continue for hours. This search has an element of the unpredictable, given that all logical deductions of the forgotten name have failed. Later, unexpectedly, the correct name often pops forth.

The possibility of subliminal searches for forgotten facts is significant. Instead of a forgotten name, one may just as well search for solutions to problems, or decisions based on past experiences. Searches with random variables occur all the time during dreaming and brainstorming. After some time, potential solutions often emerge for a person who persists in trying to find a solution. Difficult illogical problems are sometimes solved in this way, although certain problems have no solutions.

The memory model in this book attempts to calculate a *direction of attention*, a topic of immense interest not only to neuroscientists, but also to psychologists. A memory search begins with a gating of cues taken from among the features held in short-term memory. Cue subsets are selected pseudorandomly to address a variety of associated images in long-term memory. Recalled briefly and subliminally are a great many associated memories, tens per second. Such recalls alternate with impressions from sensory encoders to minimize sensory input dead time.

Each impression is flashed unnoticed against interface neurons, where a digital encoder creates an index of importance for each image. The index of importance depends on attributes in short-term memory, including bright features and strong emotions. If a subliminal recall has sufficient importance compared to what is currently in short-term memory, it is enabled by gates to become the next impression in short-term memory. This new image constitutes a direction of attention in conscious short-term memory. Direction of attention is thus achieved with a self-contained digital circuit model.

What is presented is a model of the human memory system build up from neurons. Along the way are memory searches, memorizing, and an index of importance in deciding on the direction of attention. Since the model is cast in standard analog and digital circuits, a physical structure is implied, not just an input–output formula. The result is budding cognitive architecture that may be of value for the development of artificial intelligence and intelligent robots.

General Education

This book is in its own way quite cross-curricular and strives to promote general education. General education is extremely important because it raises exponentially the number of people who are able to appreciate basic knowledge about the human memory system, thus increasing the level of support for basic research. Models are a powerful educational tool that make a subject interesting and understandable. In contrast, volumes of random, disconnected clinical facts, however academically correct, are often confusing and boring. As a way to keep students awake, the appeal of simple models cannot be denied.

CONCLUSIONS

Today, tens of thousands of researchers routinely publish what amounts to hundreds of thousands of papers every year. Seldom are these bits and pieces of knowledge reduced to a useful model. Frequently, data are published based on no model at all. The publication is justified as just another brick in a huge, overwhelming data structure. Our work here should not be considered as traditional research because it runs in a totally different direction, a direction in which existing knowledge is distilled in order to create simple models. Simple models aid education and spur interest, but they do much more.

Simple models are useful for:

1. Calculating the sensitivities of a particular neural network to its parameters
2. The discovery of some yet unknown relationship between a neural structure and its function
3. The discovery of a direction that experimental investigation ought to take

Models affect thinking

- Theories need models that are in agreement with observed facts.
- Theories with simple accurate models tend to replace theories that lack such models.
- Some theories are better than others from the point of view of model efficiency.

In this book we describe novel models of neurons, memory systems, and cognitive architectures justified in part because they point to interesting engineering applications. Modeling human memory with standard analog and digital circuits inspires computer design in two major ways:

1. A neuron may be modeled as adiabatic as far as neural signals go. This implies that the brain is an adiabatic computer: not counting calories for growth and maintenance. Energy for neural pulses is modeled as being merely borrowed; subsequently, it is returned to where it came from, the ionic solutions of the body. Adiabatic models open new possibilities for the inventive mind. For example, long-term memory circuits can be understood as neural latches with circulating signals that do not burn energy. Models like this also suggest man-made computers, electrically and logically reversible using carefully designed CMOS logic.
2. The brain is modeled as massively parallel. If everyday computers had this sort of parallelism, the world would be a very different place. Computer images would appear instantly, for example. Better yet, Web searches would be instant and accurate, actually responding only to the subject of a search. Parallelism is essential to an adiabatic computer. Indeed, without it, a CMOS adiabatic computer would be hopelessly slow.

Modeling in this book is governed in part by a principle of least action. This principle suggests that things in nature tend to operate efficiently, the underlying reason being that efficiency is necessary for species survival. Adiabatic neurons are an important aspect of efficiency. The brain, for example, performs recalls approaching tens of images per second, yet dissipates practically no detectable energy in doing so. As a tribute to the efficiency of parallel processing within the brain, millions of subconscious words of memory are instantly searched and analyzed subliminally to deliver a recall. With an eye to efficiency, we now begain to model the workings of the wonderful brain—inscrutable, but an amazement of nature with complexity beyond comprehension.

EXERCISES

1-1 A computer may be modeled in two ways: top-down, beginning with what needs to be accomplished, and bottom-up, beginning with components available for the hardware.

(a) As an example of a top-down model, sketch a block diagram of a word processor.

(b) As an example of a bottom-up model, identify hardware components to capture analog signals at the microphone input.

1-2 Compare the human brain to a man-made computer.

(a) List something for which the brain is better than the computer.

(b) List something for which the brain is not as good as the computer.

1-3 Energy and power are closely related, as a little research soon reveals.

(a) Provide an equation that relates energy to power. Define all variables and units.

(b) Provide an equation that relates power to energy. Define all variables and units.

(c) Provide conversions into joules for the following units: kilocalories, calories, electron volts, watthours, and British thermal units.

1-4 Describe an everyday situation in which parallel processing speeds up a task.

1-5 Consider grade-school arithmetic.

(a) When adding a column of integers, how can parallel operations be used? Provide a numerical example.

(b) When multiplying, how can parallel operations be used? Provide a numerical example.

(c) When dividing a smaller integer into a larger integer, can parallel operations be used? Provide a numerical example.

1-6 Based on information provided in this chapter, list the names of massively parallel associative computers from the past. Do research to describe then technically.

1-7 Name ways in which brain theory affects computer design.

1-8 Based on information provided in this chapter, sketch a block diagram of a memory system that will provide a direction of attention.

1-9 How does an imitation neuron differ from an artificial neuron as used in an artificial neural network?

1-10 List potential applications of imitation neurons.

1-11 List practical applications of artificial learning intelligence in robots.

1-12 What are the benefits of human memory system modeling?

1-13 Create, through research, a list of the hopes and fears associated with a singularity.