Alwyn Scott

How Smart is a Neuron?

A Review of Christof Koch's 'Biophysics of Computation' *

In the United States, automatic numerical computation stems from the Second World War, where the need to calculate artillery trajectories led to the development of the world's first general purpose electronic computer, called the ENIAC (for **e**lectronic **n**umerical **i**ntegrator **a**nd **c**alculator). Switched on in the summer of 1945, this machine employed 17,468 vacuum tubes connected by half a million solder joints and a tangle of plug-in cables, looking much like an enormous telephone switchboard.

Although the subsequent development of modern digital computers seems straightforward in retrospect, the path into the future was not so easy for some to appreciate at the time. As a student of physics in the early 1950s, I recall seeing an announcement of the newly invented *transistor* that included a photograph showing a small object looking like a dried pea with three wires sticking out. This spidery gob of plastic, I asked, is going to revolutionize electronics? How could such a tiny device replace the glowing output tubes of my ham radio transmitter?

A few years later, I found myself working for a semiconductor electronics company that was making transistor circuits for the computers of the ballistic missile early warning system, a giant machine designed to gather inputs from arctic radar sites and process the information to decide whether (or not) to incinerate the Soviet Union. On an idle afternoon, one of the guys speculated that it might be possible to put up to 100 transistors on a single slice of semiconductor crystal, which seemed to me a crazy idea. How could one ever check that all of the circuits were working? But the onboard computational demands of intercontinental atomic warfare (for reduced size and power consumption) soon led to the forced development of *integrated circuits*, in which thousands, then tens of thousands, then hundreds of thousands, and now *millions* of transistor circuits can be reliably interconnected on a single slice of semiconductor crystal. (Intel's Pentium III chip, for example, has 28 million

Correspondence:
Alwyn Scott, Department of Mathematics, University of Arizona, Tucson, AZ 85721, USA

* **Christof Koch,** *Biophysics of Computation: Information Processing in Single Neurons* (New York: Oxford University Press, 1999, xxiii + 562 pp., $62.95, ISBN 0 19 510491 9 (hbk.)).

transistors, corresponding to the hardware of about 1600 ENIACs, running at much higher reliability and much faster rates of data transfer.)

Going back to 1943 — when ENIAC was under secret development — an event of major significance for neuroscience occurred, and this was unrelated to the war effort (McCulloch and Pitts, 1943). Warren McCulloch, then a professor of psychiatry, and Walter Pitts, a young mathematician, published a seminal paper in which the operating principles of a digital computer were independently invented (!) and applied as a mathematical theory for the functioning of a human brain. In response to the enthusiastic promotion of this paper by physicist John Neumann in the early 1950s, researchers throughout the world began to study models of the nervous system that were based on the 'McCulloch–Pitts neuron'.

Inevitably, popular articles by science journalists about 'giant brains' began to appear in the Sunday supplements. In one sense this lurid description was literally true — for the 30 ton ENIAC filled a large room — but those starry-eyed writers intended more. In a widely circulated joke from those days, a group of engineers assemble the most powerful computer that had ever been conceived and ask it the ultimate question: 'Is there a God?' After several tense minutes of clicking and clacking and flashing of lights, a card pops out which reads: 'There is *now.*'

As is well known, McCulloch and Pitts modelled the neuron as a single switch that would fire (or not) according to whether (or not) the total input was above (or below) some threshold level. The total input was then taken to be a weighted linear sum of all of the synaptic (input) signals. Since the human brain comprises some 10 to 100 billion (thousand million) neurons, each with 1000 to 100,000 synapses, a full McCulloch–Pitts model of the brain is an enormously complex dynamic system, but is it intricate enough to model a real brain? This is the fundamental question raised by Christof Koch in his *Biophysics of Computation.*

Also appearing in the early 1950s was another classic paper by neurophysiologists Alan Hodgkin and Andrew Huxley, in which the dynamics of nerve impulse conduction along the giant axon (or outgoing fibre) of the common Atlantic squid was well described (Hodgkin and Huxley, 1952). Based upon independent measurements of parameters characterizing an isolated patch of active membrane, they developed a theory that accurately predicts both the speed and the shape of a nerve impulse, along with several other curious aspects of neural behaviour. In the context of the Hodgkin–Huxley theory, each patch of active membrane on a neuron may function like a switch, thus permitting far more intricate behaviour than is allowed to the McCulloch–Pitts model neuron.

Throughout the 1960s and early 1970s there was an uneasy truce between these two lines of research. As computers became more and more powerful, network models based on the McCulloch–Pitts neuron grew increasingly sophisticated, and corresponding engineering systems (with 'synaptic' input weights adjusted through appropriate learning routines) were employed for automatic recognition of patterns in such diverse areas as character recognition, aerial photograph analysis and cytology, among several others. During the same period, waxing computer power made it possible to analyse versions of the Hodgkin–Huxley membrane equations with ever more interesting geometries, taking account of local enlargements (or varicosities) of the fibres, branching structures in axonal (outgoing) and dendritic (incoming) trees, periodically active (myelinated) axons, and so on. Thus theoretical representations of neurons became ever more intricate.

By the early 1970s advances in electron micrography and electrophysiology made it evident to some neuroscientists that the dynamics of real neurons were far more complex than the McCulloch–Pitts model implied. Thus in 1972 neurologist Steven Waxman proposed the concept of a ‘multiplex neuron’ in which patches of membrane near the branching regions of incoming (axonal) and outgoing (dendritic) processes of the neuron are viewed as localities of low safety factor, where active nerve impulses can become extinguished (Waxman, 1972).

In other words, an individual neuron might be able to perform logical computations at its branching regions, making it more like an integrated circuit chip than a single switch. As reasonable as it seemed to some of us at the time (Scott, 1975), this expanded view of the neuron’s computational power was far from being universally accepted. Jerry Lettvin, a noted electrophysiologist at MIT, told me in the spring of 1978 that when he and his colleagues reported blockages of nerve impulses in the optic nerves of cats and speculated on the possibilities of this phenomenon for visual information processing (Chung *et al*., 1970), his NIH funding was cut off. ‘When you are ready to start doing science again,’ he was told, ‘we are ready to resume supporting you.’

Why were such seemingly reasonable suggestions so widely ignored in the 1970s? Although one can only speculate, four possibilities come to mind. (1) Admitting to increased computational power of the individual neuron would tend to undercut the validity of the extensive neural network studies that were based on the McCulloch–Pitts model. (2) Many of the studies in this area were being done in the Soviet Union (Khodorov, 1974), and Western scientists (particularly in the US) tended to ignore or disparage Soviet science. (3) The reputations of some senior scientists were invested in simpler models of the neuron. (4) There was a widespread and uncritical belief in the concept of ‘all-or-nothing’ propagation of a nerve impulse, which left no room for the concept of impulse failure at branching regions of the axonal or dendritic trees.

One of the encouraging aspects of science, however, is that the truth does out in time. Like old soldiers, the old guard withers away as evidence for the new paradigm accumulates, until what was once considered flagrant speculation becomes widely accepted as established knowledge. So it was with our collective perceptions of the neuron. Towards the end of the 1980s, as the extensive list of references in *Biophysics of Computation* shows, the experimental, theoretical and numerical evidence for impressive computational power of an individual neuron was compelling. Former heresy has become common sense.

For young scientists who are interested in understanding the dynamics of the human brain this change in collective attitude is of profound significance, to which Koch’s book provides an ideal introduction. Written in a precise yet easy style, the 21 chapters of *Biophysics of Computation* begin at the beginning, introducing the reader to elementary electrical properties of membrane patches, linear cable theory and the properties of passive dendritic trees. These introductory chapters are followed by two on the properties of synapses and the various ways that synapses can interact to perform logic on passive dendritic trees. Then the Hodgkin–Huxley formulation for impulse propagation on a single fibre is discussed in detail, and various simplifying models are presented. As a basis for the Hodgkin–Huxley description the present understanding of ionic channels is reviewed, emphasizing the importance of calcium currents. Further chapters discuss linearization of the H–H equations for small amplitude behaviour; present a careful examination of ionic diffusion processes; and describe electrochemical properties of dendritic spines, synaptic plasticity, simple neural models, stochastic neural models and the properties of bursting cells. Do you get the picture? Just about every facet of currently available neural knowledge is touched upon, with appropriate references to a carefully selected bibliography that will help the diligent novice delve deeply into whatever aspect of neural information processing he or she chooses.

All of the above comprises an extended introduction to Chapters 17 to 19, which: 'synthesize the previously learned lessons into a complete account of the events occurring in realistic dendritic trees with all of their attendant nonlinearities'. 'We will see', the author writes, 'that dendrites can indeed be very powerful, nontraditional computational devices, implementing a number of continuous operations.'

Thus *Biophysics of Computation* offers a definitive statement for the direction in which the neural research of the new century should go. Chapter 20, the penultimate, discusses several speculations for non-neural computation in the brain, ranging from molecular computing below the level of a single neuron to the effects of chemical diffusants (nitric oxide, calcium ions, carbon monoxide, etc.) on large numbers of neurons. Although this entire area has been neglected by most of the neuroscience community, Koch points out that there are no good reasons for doing so. As we enter the new century, neuroscientists should keep their minds open.

Finally, in the summary of Chapter 21, seven problems for future research projects are listed, emphasizing that the investigation of information processing in single neurons is very much a work in progress. It is of interest to examine these 'strategic questions' as they reveal the author's intuitions about possible directions of future developments. (Note that these are not direct quotes, as I have taken the liberty of summarizing Koch's questions.)

(1) How can the operation of multiplication be implemented at the level of a single neuron?
(2) What are the sources of noise in a neural system and how does this noise influence the logical operation of a single neuron?
(3) How is the style of neural computation influenced by metabolic considerations?
(4) What is the function of the apical dendrite, which is a typical cortical structure?
(5) How and where does learning actually take place in a neural system?
(6) What are the functions of the dendritic trees, the forms of which vary so widely from neuron to neuron?
(7) How can we construct neural models that are sufficiently realistic to capture the essential functions of real neurons yet simple enough to allow large-scale computations of brain dynamics?

As these questions indicate, Koch is not merely concerned with understanding what unusual behaviours the neuron does or might exhibit. His broad aim is to comprehend the relation between this behavioural ability and the computational tasks that the neuron is called upon to perform. In his words:

> Thinking about brain style computation requires a certain frame of mind, related to but distinctly different from that of the biophysicist. For instance, how should we think of a

chemical synapse? In terms of complicated pre- and post-synaptic elements? Ionic channels? Calcium binding proteins? Or as a non-reciprocal and stochastic switching device that transmits a binary signal rapidly between two neurons and remembers its history of usage? The answer is that we must be concerned with both aspects, with biophysics as well as computation.

This excellent book is evidently a labour of love, stemming from the author's 1982 doctoral thesis on information processing in dendritic trees. As far as I can tell all relevant aspects of neural processing are considered, with what seem to me to be just the proper amounts of emphasis. The writing style is precise and rigorous without being stuffy, and the many references to a fifty-page bibliography will be of enormous value to young researchers starting out in this field.

In addition to its obvious value for those engaged in experimental, theoretical or numerical studies of neuronal behaviour *Biophysics of Computation* would also work well as the text for an introductory course in neural dynamics, perhaps as part of a neuroscience programme. Considering the rate at which new experimental information is accumulating in this area (from the ongoing efforts of myriad electrophysiologists and electron microscopists), and taking account of the ever increasing availability of computational power, it seems inevitable that computational biophysics will become established as a major research area within the next few years. Those who would enter this exciting arena must make themselves familiar with the contents of Koch's book.

As is often the case, the writer of this positive review feels compelled to offer a few reservations, if for no other reason than to demonstrate that he is unbiased and has read the entire book. Indeed Koch invites such comments for consideration in future editions. Mine are twofold.

First, I was disappointed to find no mention of Steve Waxman's 1972 paper in which the concept of the multiplex neuron was introduced (Waxman, 1972). Although largely ignored by the neuroscience community of the early 1970s, this publication proposes the same broad view of neural intricacy as does *Biophysics of Computation.* Certainly Koch's list of references needs to be carefully selected, and he has done so, but Waxman's seminal work deserves a mention.

Second, there are very few references to the many significant publications of Soviet scientists on the dynamics of individual neurons, especially those appearing during the 1970s (Khodorov, 1974). This is an unfortunate oversight that could be rectified in a future edition by asking for suggestions from some of the leading Russian neuroscientists.

Having established that *Biophysics of Computation* is an important addition to the literature of neuronal information processing, one question remains: Why should readers of *the Journal of Consciousness Studies* find it interesting?

To answer this question, consider the enormous influence that the McCulloch–Pitts paper has had on the cognitive science community since its discovery and promotion by Neumann in the late 1940s. This subtle influence is not, I believe, unrelated to the subsequent journalistic anthropomorphizing of ENIAC and her increasingly sophisticated offspring. From this early perspective each neuron is regarded as a single switch, and the human brain is likened to a digital computer, as we currently understand the meaning of that term. This limited model of the brain has been widely disseminated throughout the realms of science, leading many to conclude that the human brain *is* essentially a digital computer. Others suggest that mysterious physical phenomena—such as quantum theory—must be brought in to neuroscience to square what we know about the nature of mind with the facts of neurology.

From the perspective of Christof Koch's *Biophysics of Computation* the situation is quite different. A neuron can no longer be viewed as a single switch; it is more or less analogous to an integrated circuit chip. I write 'more or less' because much of a neuron's behaviour remains under a shroud, making it difficult to discern what a system of 10 to 100 billion real neurons, as described in this book, might or might not be able to do. Cognitive scientists should give priority to accepting and appreciating these emerging vistas and offering re-evaluations of the brain's abilities for the benefit of those psychologists, philosophers and physicists who are engaged in consciousness studies.

References

Chung, S.H., Raymond, S.A. and Lettvin, J.Y. (1970), 'Multiple meaning in single visual units', *Brain, Behavior, Evol.*, **3**, pp. 72–101.

Hodgkin, A.L. and Huxley, A.F. (1952), 'A quantitative description of membrane current and its application to conduction and excitation in nerve', *J. Physiol.*, **117**, pp. 500–44.

Khodorov, B.I. (1974), *The Problem of Excitability* (New York: Plenum Press).

McCulloch, W.S. and Pitts, W.H. (1943), 'A logical calculus of the ideas immanent in nervous activity', *Bull. Math. Biophys.*, **5**, pp. 115–33.

Scott, A.C. (1975), 'The electrophysics of a nerve fiber', *Rev. Mod. Phys.*, **47**, pp. 487–533.

Waxman, S.G. (1972), 'Regional differentiation of the axon: A review with special reference to the concept of the multiplex neuron', *Brain Res.*, **47**, pp. 269–88.