6

Group Report: Molecular and Biophysical Mechanisms of Information Processing

C. KOCH, Rapporteur

H. BETZ, R.J. DOUGLAS, M. EGELHAAF,
J. JACK, H. KORN, E. MARDER, H. PARNAS,
L.A. SEGAL, R.D. TRAUB, T. VETTER

INTRODUCTION

"By their fruits ye shall know them" —Matthew 7:20

The brain computes, that is, it transforms immense amounts of highly redundant sensory information into decisions—taking into account various internal states—that ultimately manifest themselves in motor output. But how do these computations occur? In particular, how are the constitutive elements of nervous systems, the neurons, involved in such computation? And to what extent can they be said to instantiate computations? Furthermore, how can models of the nervous system, or parts of it, be evaluated in terms of such computations expressed at the neuronal level? The task of our group was to evaluate these and similar issues.

Before launching into this task, we would like to emphasize that models at the biophysical and network levels have found widespread acceptance within the neuroscience community. Modeling as a formal, quantitative means of describing neuronal systems has become a powerful tool in the repertoire we have available for understanding the nervous system. This may reflect the shift of neuroscience from a descriptive science to a more explanatory, physics-like science. The growth of modeling is due to a number of distinct causes:

1. *Limitations of Intuition*: Modeling transcends the limitations of human intuition when faced with complex neuronal systems. Neurobiology, in particular as practiced at the biophysical or circuit level, contains a large number of highly

nonlinear elements, e.g., ionic channels, bursting or oscillating neurons, which severely limit the ability of the researcher to understand the system and make predictions. Extreme sensitivity to parameters and to the cascades of nonlinear elements preclude any simple analysis and leave quantitative models as the only alternative. This situation is quite different from that of fifteen, ten, or even five years ago, where modeling, together with computational considerations, was marginalized and met with considerable scepticism.

2. *Data*: Detailed and comprehensive data concerning the kinetics and distribution of ionic channels and neurotransmitters, the morphology and distribution of nerve cells and their connectivity patterns, and how these change over time are becoming available. Data of this kind are best organized in the form of functional models that can be used to test the consistency of the data and draw attention to areas that require further research.

3. *Hardware*: The widespread availability of massive and cheap computing power, mainly in the form of easy-to-use, single-user, UNIX type workstations with graphical abilities, has enabled more researchers to participate in modeling directly. Supercomputers such as the CRAY Y–MP or the Connection Machine are used mainly for detailed investigations of network dynamics.

4. *Mathematical Tools*: It has been fortuitous that the formalism introduced by Hodgkin and Huxley in 1952, in which ionic conductances are modeled in terms of one or more fictive activation and inactivation "particles" with first-order kinetics, has proven adequate to model, with various modifications, the ever-growing number of ionic currents, including calcium-dependent ones. This formalism, together with cable theory, represents the framework within which single cells and networks have been analyzed.

Neural modeling efforts[1] have led to a number of successful predictions. For instance, the theoretical analysis of Rall and Shepherd (1968) of the laminar field potential in the olfactory bulb predicted the existence of (reciprocal) dendro-dendritic synapses between individual granule and mitral cells. These synapses were independently observed in an electronmicroscopic study by Reese (Rall et al. 1966). Another instance is the model of synchronization of neuronal discharge in the disinhibited hippocampal slice by Traub and Wong (1982). It predicted that bursts in a single pyramidal cell could evoke bursts in one other pyramidal cell (confirmed with the aid of dual recording by Miles and Wong 1986) and also that stimulation of a single neuron could evoke synchronous firing in a large (e.g., thousands of cells) population of disinhibited cells (experimentally verified by Miles and Wong 1983). Another prediction comes from the work of Kepler, Marder and Abbott (1990) who studied a circuit consisting of a single oscillator electrically coupled to a nonoscillatory neuron.

---

[1] Here we only consider neuronal models, i.e., models which ultimately express high-level behavior, perception, action, etc. in terms of the interaction among many neuron-like components; this rules out more abstract theories, such as the black-box models of cognitive science.

This work was motivated by the experiments of Hooper and Marder (1987), who showed that the frequency of the bursts generated by the anterior burster (AB) neuron was decreased by its coupling to nonbursting cells. This model made certain predictions regarding the interactions of oscillatory and nonoscillatory neurons that were later confirmed experimentally (Sharp et al. 1992).

Models in neuroscience can be classified into two broad, nonexclusive categories (see also Sejnowski et al. 1988): One group of models has as its principal aim the detailed and quantitative description of the system under study. Knowledge of the individual components, usually the behavior of single cells or voltage-clamp data from individual ionic currents, is used to predict the evolution of the entire system in time. Such an analysis helps to pinpoint the shortcomings of the original data, highlights the key role of certain parameters and, if successful, reassures one that the behavior under scrutiny can indeed be explained within the current framework. Examples of such models span the entire gamut of spatio-temporal scales, from the analysis of quantal release at the afferent synapse to the Mauthner cell (Korn and Faber 1992) to firing properties of sympathetic ganglion cells (Yamada et al. 1989) and cortical pyramidal cells (Stratford et al. 1989; Koch et al. 1990; Lytton and Sejnowski 1991), to the modeling of small networks of bursty cells in the lobster (Marder 1992, Marder et al. 1992) and the large-scale modeling of hippocampus, olfactory and visual cortex (Traub and Miles 1991; Wilson and Bower 1989; Bush and Douglas 1991; Wörgötter and Koch 1991). Note that such models can also be considered as a compact description of the properties essential for the behavior of the system.

In contrast to these models motivated primarily by biophysics, computational models seek to study how particular information processing tasks, e.g., computing the direction of motion or depth from binocular disparity, can be instantiated in neuronal hardware. Here the emphasis is on the computational metaphor with a concomitant reduction of biophysical or anatomical detail. Such models come in very different flavors of neurobiological realism, from the investigation of the interaction between excitation and silent inhibition in retinal ganglion cells (Koch et al. 1982) and the analysis of the circuits underlying the detection of relative motion in the fly (figure-ground discrimination; Reichardt et al. 1983; Egelhaaf et al. 1988) to the conventional connectionist or Hopfield (1984) class of neural networks (as well as their intellectual spawn) for perceptual, motor, and memory tasks. In each case, the emphasis is on the processing (e.g., AND–NOT gates in the Koch et al. 1982 study) and the representation of the information (unit vs. place coding in connectionist systems; Ballard et al. 1983).

Of course, today many models cannot be clearly segregated into one or the other category and represent an intermediate type. However, the primary intent behind the constuction of the model is usually either a descriptive or a computational one.

Over the last five years or so, a new simulation, or more correctly, emulation, technique has emerged: analog very large-scale integrated circuits (VLSI) fabricated using CMOS design methodologies and technology. This approach towards building neuronal circuits has been pioneered by Mead (1989) at Caltech where he and his group have constructed a number of neurobiological circuits, ranging from the silicon

retina (Mead and Mahowald 1988) and cochlea (Lazzaro and Mead 1989) to circuits mimicking the vestibulo-ocular reflex (DeWeeth and Mead 1990) and central pattern generators in various invertebrates (Ryckebusch, Bower and Mead 1989). More recently, Mahowald and Douglas (1991) have successfully built an analog VLSI model of a cortical pyramidal cell, including a number of adapting currents (I_M, I_{AHP}, and I_A). The key difference between these circuit emulations on the one hand and digital computer simulations or digital hardware on the other is that the device physics is exploited to implement the various primitives found in biological circuits. For instance, the current flowing through a field effect transistor biased in the subthreshold range is exponential in the gate voltage, with an effective "gating charge" of 0.7. Such a circuit element can be used to mimick the activation or inactivation variable associated with voltage-dependent conductances (Mahowald and Douglas 1991). Mead (1989) argues that the physical constraints imposed on the density of wires in the essential two-dimensional world of electronic circuits as well as the noisy nature of these analog components are very similar to the type of constraints imposed on biological circuits. These methods are radical different from computer simulation, because they force the VLSI designer to implement very robust and adaptive algorithms with an eye towards minimizing unnecessary communication among elements. As silicon foundry services (e.g., MOSIS) and VLSI design tools become more widely available, this simulation technology will become more widely used.

FROM BIOPHYSICS TO COMPUTATION AND BACK

What is meant by the term "computation" or by the phrase "this neuron implements a multiplication"? In the most general sense a computation represents a fixed mapping or transformation from the states of one physical system to the states of another physical system, with the mapping corresponding to a given function. Examples include the position of two marked pieces of wood implementing multiplication (exploiting the fact that $\log(xy) = \log(x) + \log(y)$), or a marble rolling downhill and coming to rest at the bottom of the valley, implementing the minimization of a two-dimensional functional via the use of gravity, or an electrical charge flowing onto digital gates, implementing a digital NAND, or the action of a transient, inactivating potassium current (I_A) implementing a temporal delay in certain neuronal systems. In order to instantiate a sufficiently powerful class of computations, in particular in order to make decisions, certain minimalist requirements need to be met, most importantly the existence of some nonlinear element.

These requirements, however, are very weak, as argued by D. Hillis as well as R. Rivest. It is obvious that general models of computation (i.e., simply the requirement that the system must be able to "compute") do not provide enough constraints to help understand computations in neuronal systems. What is needed to build a computational engine, from a practical point of view, are units with a nonlinear input–output function and some memory units, a mechanism for the amplification and distribution

of signals to other units (a high fan-in and fan-out) as well as a means of fast communication among units. However, the detailed nature of the elementary computing units, e.g., whether they implement a threshold function or a squaring unit or whether they are nonstationary, does not matter a great deal. Thus, since almost any model of a neuron fits the above requirements, these constraints are too weak to truly help us.

One reason why existing computational models seem to tell us so little about the design of biophysical circuitry is that the data structures we use in our computational models do not map in a natural way onto the biophysical medium. The assumption made in most modeling efforts is that details of the biology can be ignored when specifying the function to be computed and the algorithm and data structures to be employed, because the biophysical medium is powerful enough to implement whatever functions we have independent reason to believe are being computed. As D. Kirsh pointed out, the issue highlighting this problem is our ignorance about possible models of computation. At present we operate with a small set of formal models of computation, such as Turing machines, lambda calculus, and Boolean networks. Traditional computer science is built on the assumption that data structures such as graphs, balanced trees etc. will work for formal analyses of computations. These are all pointer-based notions and may prove to be wrong or irrelevant data structures for biological systems where timing, connectivity, etc. is so important.

All these problems do not invalidate the usefulness of a more detailed computational analysis in the sense of the word used by Marr and Poggio (1977). For instance, an analysis of the motion problem at the abstract computational level leads to the realization that estimation of the two-dimensional optical flow from the time-varying intensities falling onto the retina is an ill-posed problem requiring the use of a smoothing stage. A similar analysis of the problem of extracting depth from two-dimensional images leads to the insight that a number of constraints are needed by *any* system, whether a fly, a cat, a Californian surfer, a robot or an analog VLSI chip, to solve this problem (smoothness, forbidden zone, figural continuity, spatial coincidences, epipolar lines, etc.). Since they do not depend on any particular algorithm or hardware, such constraints are exceedingly useful in attempting to understand how a particular system solves this problem. Thus, we reaffirm the notion that such constraints are very useful and are probably a *conditio sine qua non* for understanding complex information processing systems.

Furthermore, we all believe to a greater or lesser extent that the inverse route, from biophysics and circuits to computation, is both possible and desirable, in the sense that we will ultimately have some understanding of which biophysical mechanisms or circuits are responsible for which computational task. Where we differed substantially within our group was how much "blood, sweat, and tears" this program is likely to cost and what its limitations are likely to be. For instance, in many systems of tightly interlocking elements it may be very difficult to isolate the contribution that any individual component makes to the overall task or a synergistic effect may be involved, with many mechanisms contributing simultaneously.

C. Koch et al.

MODELS OF COMPUTATION IN SINGLE CELLS

The vast majority of neural network models contain few if any of the plethora of biophysical and anatomical complexities we know are present in real neurons. At the other extreme, there exist models of Purkinje cells with a thousand or more compartments, each one representing, in great detail, a small fraction of the cell and its synaptic input. Along the middle ground would be a "canonical" or "simplified" model of a neuron which, while more faithful to biological reality, does not overwhelm us with its wealth of—possibly irrelevant—detail. Conceptually, such a neuron can be subdivided into a small number of functional compartments:

1. the postsynaptic region, translating local transmitter input into a conductance change;
2. one or several dendritic compartments performing spatio-temporal integration of the membrane potential;
3. a response generator converting the current flowing from the dendrite into the cell body into a pattern of action potentials;
4. a nerve terminal transducing the action potential into probabilistic release of packets of neurotransmitter.

The mechanism of release of neurotransmitter from nerve terminals has been described by the collective work of Katz, del Castillo and Miledi (for a summary, see Katz 1966). According to their work, the action potential arriving at the nerve terminal causes membrane depolarization, resulting in the influx of Ca^{2+}. Neurotransmitters are stored at the synaptic terminal in synaptic vesicles, each of them containing about 4,000–10,000 molecules of a given mediator, e.g., ACh or glutamate. These vesicles are released from specialized (active) zones following the calcium influx and bind to postsynaptic receptors. The release of vesicles is not deterministic but is governed by probabilistic laws (Katz 1966).

One essential aspect of transmitter release concerns the probabilistic or stochastic nature of neurotransmitter release at central neurons (see Korn and Faber, chapter 2; Redman 1990; Korn and Faber 1987; Korn et al. 1981; Larkman et al. 1991; Jack et al. 1990). In the case of central neurons, one anatomical synapse usually contains one active release zone. Because a given neuron makes many synapses onto a given postsynaptic target in its terminal arborization, the number of active zones, n, for a given connection among two neurons ranges from as few as 1 or 2 to as many as 100, with $n = 5$–15 being the most common for inhibitory synapses. Vesicles are released independently at each active zone, either spontaneously with a very low rate of occurrence (less than 0.1 Hz) or synchronously at some of the terminals, following an impulse. It is now generally accepted that a single, rather than many, vesicle is released at a single active zone; the probability of this all-or-none release is p. Besides these two parameters, n and p, a third parameter q refers to the size of the postsynaptic effect produced at the postsynaptic location by the binding of the released neurotransmitter

and the subsequent opening of ionic channels. q can be thought of as the unitary postsynaptic conductance change following the release of one packet of neurotransmitter. The probability distribution associated with the vesicle release approximates a binominal distribution; either a simple one, with p the same for all active zones, or a compound distribution, where p differs from one active zone to the next. This general scheme varies in its details for different cells.(Korn and Faber 1992); however, these three parameters, n, p, and q, characterize the process and must therefore in many cases replace the simplistic assumption of a fixed, deterministic synaptic weight T_{ij} (or w_{ij}) between neuron i and neuron j in neural network models. The possible consequences of the stochastic release for neuronal computations have been explored by Burnod and Korn (1989).

It is important to realize that the parameter q depends to some extent on the postsynaptic targets (Korn and Faber 1991). Also, p can vary with firing frequency, from very low values ($p = 0.05$–0.1) to higher ones (0.3–0.6). This feature manifests itself as plasticity, with facilitation in the case of low release probability and depression in the case of large p values. Furthermore, p appears to be inversely proportional to n, so that the mean release, np, remains invariant with changes in the number of synapses. The postsynaptic quantal response, q, also appears to vary with properties of the postsynaptic cell; in the case of the motoneuron a quanta on a distal dendrite may open up to 5–10 times more channels than a quanta on the soma (Jack et al. 1990).

Considerable progress has been made in recent years in formulating and analyzing quantitative models of synaptic release. It is agreed that the membrane depolarization following invasion of the presynaptic terminal by an action potential opens calcium channels in the membrane, elevating intracellular calcium. This calcium then brings about the exocytosis of the vesicles containing the neurotransmitter. One group of models (Fogelson and Zucker 1985; Simon and Llinas 1985) asserts that the spatio-temporal patterning of intracellular calcium alone is the necessary and sufficient condition for vesicular release. The alternative calcium-voltage hypothesis (Parnas et al. 1986) is based on the assumption that calcium is essential but not sufficient in itself to evoke release of neurotransmitter. According to this detailed model—formalized in compact form by Lustig et al. (1989)—an additional molecule must be activated by membrane depolarization which binds calcium. In its bound form this molecule induces exocytosis.

A feature neglected by the majority of the modeling community is the ability of neurons to respond in a qualitatively very different manner to synaptic input activity (McCormick 1990). A tonically active cell fires one or more action potentials, in a manner reminiscent of the Hodgkin–Huxley squid axon. Increasing the input current leads to an increase in spiking activity. Bursting cells discharge by generating fast sodium spikes superimposed onto a broader calcium-mediated response. Following such a burst, the membrane is hyperpolarized for some time (between 20–200 msec) and remains silent. After a while, the cell may once more begin to fire a burst and the cycle begins anew. Thus, the patterns of action potential discharge in the tonically firing versus the bursty cell are quite distinct. In cortex, such bursting cells appear to

be predominately pyramidal cells localized to layer V. Why cells discharge in these two patterns is at present not known. Finally, a third type of cellular response consists of a plateau potential, that is an equilibrium state in which the membrane is depolarized and may discharge action potential (Marder 1992). Such cells behave in a flip-flop manner: once stimulated they will remain in their firing state until the membrane is sufficiently hyperpolarized by an external current input. At that point in time, they will return to their stable quiescent state.

As described in the background paper by Marder (1992), input by various neuro-modulators can switch one cell from a bursty firing mode into a tonic firing mode. Such behavior has been observed in both invertebrate ganglions (Marder 1992) as well as in the mammalian thalamus (McCormick and Pape 1990; Pape and McCormick 1989) and appears to correlate in the latter case with the gross behavioral categories of deep sleep and wakefulness. Since cells respond radically differently in these states, it is crucial that models incorporate this behavior.

Besides engaging in state-dependent switching, neuromodulators such as noradre-naline, acetylcholine acting on the muscarine receptors, serotonin, and the ever-grow-ing list of neuropeptides can also vary the spatial and temporal time scale over which neurons can act. This diversity of actions is enhanced at the molecular level by cell-specific regulation of receptors and their gene expression as well as protein modification, giving rise to an enormous substrate of mechanisms possibly regulating synaptic efficacy and plasticity (Betz 1992).

In the hippocampus, activation of noradrenergic fibers leads, via a second messenger, to the blockage of one or more potassium currents, thus controlling spike frequency adaptation and consequently the gain of the neuron. Likewise, LHRH-like peptide acting on the bullfrog sympathetic ganglion cell can increase the gain of the neuron over many seconds (Adams et al. 1986). The dominant feature of most of these neuromodulators is their temporal time scale. Their duration is in the second to minute range.

It is important that realistic models incorporate this diversity of time scales: while the passive membrane time constant in a quiescent cell may be in the 10–20 msec range, it can drop to 3–5 msec during action potential discharge. Firing frequency adaptation is controlled over tens of msec to hundreds of msec, while the effects of modulatory substances extends for seconds and minutes.

The dominant model of the dendritic tree is the one developed by Rall (1989). This model assumes that synaptic inputs can be treated as current inputs and that the membrane does not contain any (significant) membrane nonlinearities. The resultant theory describes reasonably well the response of alpha motoneurons to one or a few synaptic inputs from Ia afferents (Burke 1990). A number of other cells, such as subclasses of retinal ganglion cells, also appear to possess passive dendritic trees, devoid of any significant membrane nonlinearities. However, the majority of neurons do contain significant membrane nonlinearities in the dendritic tree. In particular, at least two classes of calcium channels have been observed in the dendritic tree of most central neurons, as have been fast and sustained sodium channels and a number of potassium channels (Llinas 1988; McCormick 1990).

The majority of modelers have neglected voltage-dependent conductances in the dendrites (for an exception see Traub et al. 1991), both for reasons of computational complexity as well as for paucity of hard data concerning the detailed distribution of these currents and their kinetic description. This lack of data is grounded in the difficulty of directly recording in the dendritic tree. Consequently, conclusions concerning the nature of active conductances usually have to be inferred from recordings done at the cell body. This biophysical complexity could, at least in principle, give rise to computational complexity, in which, for instance, dendritic spikes subserve threshold computations at many different dendritic locations prior to the spiking threshold at the soma (Llinas and Sugimori 1980; Mel 1992).

A further complexity is the presence of voltage-dependent synaptic conductances, in particular the glutamate-activated n-methyl-d-aspartate (NMDA) receptor. The associated channel displays voltage-dependent behavior due to the blockage of the channel by Mg^{2+} at voltages around the resting potential (i.e., -70 mV); more depolarized potentials relieve this block. Since this ionic channel requires a conjunction of presynaptic neurotransmitter and postsynaptic voltage depolarization in order to open, it has been proposed to underlie Hebbian learning. Indeed, experimental evidence implicates NMDA receptors in the induction of some forms of associative LTP in the hippocampus (Bekkers and Stevens 1990; Brown et al. 1990). Detailed biophysical modeling (Mel 1992) has shown that a cell endowed with a population of mixed NMDA and non-NMDA synapses approximates a multiplicative-like behavior, instantiating a so-called sigma-pi neuron (in such a formal model, synaptic inputs interact in a multiplicative-like manner prior to summation at the soma). Poggio and Girosi (chapter 7) argue that such units could underlie interpolation-like learning schemes (look-up table learning).

Following a discussion of these and other biophysical mechanisms possibly underlying information processing (such as dendritic spines, the influence of synaptic location on function, gap junctions, dendro-dendritic synapses, branch point failures, etc.), we agreed on "a" simple neuronal model (the *Dahlem* neuron) which, while reflecting some of the biophysical complexity, would not impose too large a burden on researchers studying the behavior of large neuronal networks. Such a simplified neuron should include:

1. elementary, time-dependent, synaptic-induced conductance changes with variable time constants;
2. multiple overlapping time scales in the conductances controlling excitability;
3. fast Hodgkin–Huxley-like spiking mechanism;
4. bursts and plateau potentials (if they exist in the cell type investigated);
5. a small number (2–4) of dendritic compartments;
6. topology of the synaptic network interconnectivity.

What is important in this list is not the detailed form of, say, the conductance changes or of the fast spiking mechanism, but that these mechanisms are included.

Koch and Poggio (1987) proposed that a short list of biophysical mechanisms and the neuronal operations they implement would prove to be very useful in understanding neuronal information processing. This approach is similar to C. Mead's (Mead and Conway 1980), who discussed the small number of physical effects underlying computations in digital computer systems. Koch and Poggio argued for a similar "biophysics of computation" program, in which various biophysical mechanisms are thought of as implementing specific formal operations, such a gain control, multiplication, linear addition, etc. Their list has been recast and elaborated by Douglas and Parnas in Fig. 6.1, which emphasizes the direct mapping of some significant biophysical mechanisms and their neuronal operations onto a model neuron. In some sense, it is our hope that this *Dahlem* neuron will replace the much, much simpler McCullough and Pitts neuron as the primitive building block in neuronal networks. It is important to point out that the mechanisms listed are a subset of the mechanisms known to exist in nerve cells. In particular, none of the biochemical or biophysical mechanisms modulating the neuronal response via second messenger systems and phosphorylation of ionic channels have been included here. Their time course makes them less relevant for "fast" neuronal computations of the kind underlying perceptual or motor tasks (i.e., a brief fraction of a second), but crucial for subserving adapation, plasticity and learning.

MODELS OF COMPUTATION IN NETWORKS

The style of modeling neural networks is a trade-off between biological verisimilitude, computational load, and functional abstraction. Thus the models range from very detailed "rococo" models, with the number of neurons and types matching those found in a particular biological circuit, to "baroque" simplified neurons in circuits that preserve average connectivity, to the "postmodern" and "minimalist" networks composed of highly simplified neural elements and the austere connectivity of artificial neural networks.

"Rococo" Models

Typically the number of "neurons" N_{model} in these simulations satisfy $N_{model} \cong O(N_{biology})$, with an accurate connection topology (Traub and Miles 1991; Wilson and Bower 1989). The individual neurons are usually of the compartmental type, with the detailed dendritic geometry reduced to a few (4–20) compartments. The somatic compartment contains active conductances that replicate the intrinsic excitability characteristics of the particular cell types; active currents are added to dendrites *ad libitum*. The conductances follow Hodgkin–Huxley-like dynamics, and the output will demonstrate tonic, bursting, and adaptive discharges. Lack of experimental data may limit the accuracy of the kinetics; however, both the membrane as well as the synaptic currents should have correct overlapping of time scales. The fundamental idea
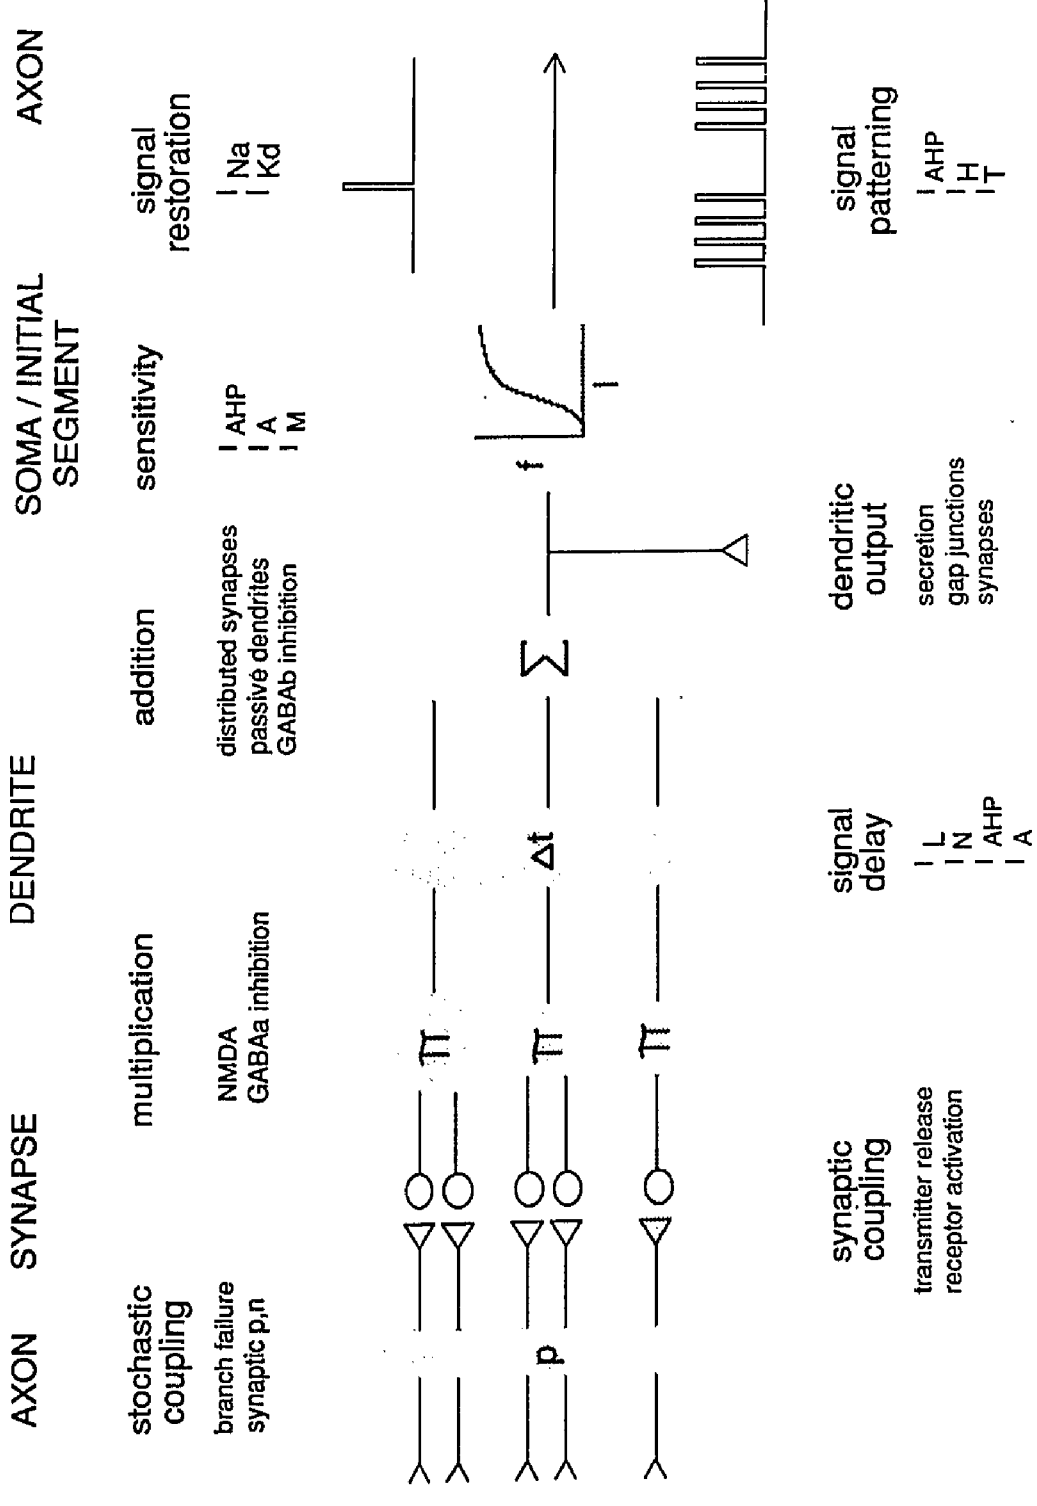
Figure 6.1 The *"Dahlem"* formal neuron represents an attempt to map some of the known biophysical mechanisms and ionic currents onto more abstract neuronal computations (see also Koch and Poggio 1987). While this diagram is by no means exhaustive, it does include most of the important components thought to play a role in most cells throughout the animal kingdom. Note that the formal operations, such as signal restoration, multiplication etc., are approximations to the more complicated events occuring at the membrane level.

underlying this sort of modeling is to create as accurate a model of the biological circuitry as possible, with the aim of comparison with physiological experiments.

"Baroque" Models

In these models, simplification is achieved by replacing detailed connectivity with average connectivity and by considering the interaction of populations of neurons rather than individual neurons (see the target article by Douglas and Martin 1992). Each population is represented by a detailed neuron with characteristics similar to those described in the previous section. However, the spike output is encoded in terms of firing frequency and used to represent the average discharge rate of the population. The statistical variability of the population discharge can be approximated by convolving the spike output of the reference neuron with a suitable interspike interval distribution. Models of this kind have been used to investigate the synchronization of discharge in neocortical pyramidal cells (Bush and Douglas 1991).

"Postmodern" Models

Neurons in such models usually have two state variables: a fast variable corresponding to voltage and a second, slower variable corresponding to accomodation or adaptation. The best known one is the singularly perturbed system of FitzHugh (1961) and Nagumo et al. (1962). These models reproduce most of the dynamical behavior seen in biophysically more faithful models, but at a much reduced complexity, and have lead to novel—and experimentally verified—predictions (Guttman et al. 1980). These single cell models have also been incorporated into simple networks and their temporal evolution has been studied (Rinzel and Ermentrout 1989; Abbott 1990).

"Minimalist" Models

These are typical artificial neural network models (Hopfield 1984; McClelland and Rumelhart 1988). The neurons are reduced to "neural-elements" that are usually linear-threshold elements with graded input and output. They are connected in very simple topologies, for example, three-layered, feedforward networks. These networks are used to investigate synaptic learning and pattern recognition tasks or for analyzing various perceptual tasks in cortex, such as optical flow or binocular disparity. Despite their simplicity they are able to solve some interesting problems (Lehky and Sejnowski 1988; Zipser and Andersen 1988; Wörgötter and Koch 1991).

The type of model that will be used for a particular application is a matter of intuition and degree of experimental detail available as well as computational cost. It is instructive to consider how properties of neuronal networks scale with size.

As argued by Traub, the number, I, of synaptic inputs to a particular cell depends on N, the number of neurons in the population, and p, the probability of connection, such that $I = Np$. Clearly, if I and p must be preserved, then the number of cells in the

model must agree with that of the original network. If it is necessary to reduce the number of cells in the model so that $N_{model} \ll N_{biology}$, then it will also be necessary to change I, p, or both. Which strategy is used will depend on the problem in question. For example, in an investigation of epileptic synchronization, it is the chain reaction of excitation between neurons that is important. In this case, connectivity is irrelevant. The model would be effective provided that $I > 1$, which ensures at least one connection between cells (Traub and Miles 1991). In other cases the nature of the physiological process may preclude reduction of N. Suppose we wish to understand how recurrent inhibition regulates the influence that one pyramidal cell exerts on another, randomly chosen, pyramidal neuron (Miles and Wong 1987). It turns out that the excitatory effect of one neuron percolates along a series of postsynaptic targets, progressively increasing the total network excitation. Clearly this effect depends on the propagation of excitation across the network of individual neurons, and correct simulation requires that $N_{model} \cong O(N_{biology})$.

Figure 6.1 shows that individual neurons can be decomposed into biophysical mechanisms, which can be interpreted as computational "operations." This interpretation is influenced by a digital/logic view of neuronal computation. Moreover, it is restricted to the level of single neurons. A different approach emphasizes analog signal processing at subneuronal, neuronal, and network level (Shepherd 1972; Mead 1989; Douglas et al., in preparation). By contrast with Fig. 6.1, this view stresses the adaptive processes that adjust the performance of neurons to match their input, gain control, signal/noise enhancement, signal restoration between computational stages, short- and long-term memory, and signal averaging. The multi-level and multi-timescale nature of these processes is exemplified by adaptation. In this case, the adjustment occurs on a number of time scales and acts, e.g., via up-down regulation of channels, the balance of activation of synaptic vs. extrasynaptic inhibitory channels, short- and long-term constant feedback inhibition, and changes in dendritic time constant and length constant due to prevailing spontaneous synaptic input. Rules of this type will then allow us to reduce complicated circuits into simpler signal processing modules.

CONCLUSION

1. Neuronal modeling at both the detailed biophysical/anatomical level as well at the more abstract computational level is widely accepted in the neuroscience community today and represents an important tool in the armory of the neurobiologist.

2. Given the immense diversity in receptors, ionic channels, and neuronal morphology observed across species and within species in different parts of their nervous system, each particular neuron and circuit needs to be modeled in its own right. We did, however, generate a figure, formalized as the *Dahlem* neuron, that emphasizes the more important biophysical mechanisms underlying neuronal computation. Thus, it appears that the autocratic reign of the

C. Koch et al.

simplified McCullough-Pitts neuron is being replaced by a democratic rag-bag of neurons with very eclectic properties. Variations from one neuron to the next are grounded in the evolutionary nature of the system and we may not be able to derive them from computational considerations.

3. We simply do not yet know to what extent the observed biophysical and anatomical complexity is mirrored in computational complexity.[2]

4. We are agnostic when it comes to making the transition from the simple systems we can model today to the quantitatively much larger brain circuits we would like to understand. As argued by Marder, our collective experience with circuits of two or more neurons in invertebrates has shown that intuition is overwhelmed by our inability to predict the consequence of small variations in various parameters in these strongly coupled and highly nonlinear feedback systems. This is one of the principal reasons experimentalists have enthusiastically endorsed and participated in neuronal modeling. At the moment, only further trial-and-error experimentation and the further development of analytical models of neuronal interactions will enable us to tell to what extent the chiasm between simple neuronal systems and large-scale neuronal networks can be bridged.

5. This calls for novel experimental techniques. Thus far, electrophysiologists have been inhibited in their approach to large networks by the embarrassing smallness of their tools (Douglas and Martin 1991). It is crucial that larger tools be employed to derive insight into activities occurring in large neuronal structures.

REFERENCES

Abbott, L.F. 1991. Realistic synaptic inputs in network models. *Network: Computation in Neural Systems* 2:245–248.

Adams, P.R., S.W. Jones, P. Pennefather, D.A. Brown, C. Koch, and B. Lancaster. 1986. Slow synaptic transmission in frog sympathetic ganglia. *J. Exp. Biol.* 124:259–285.

Ballard, D.H., G.E. Hinton, and T. Sejnowski. 1983. Parallel visual computation. *Nature* 306:21–26.

Bekkers, J.M., and C.F. Stevens. 1990. Computational implications of NMDA receptor channels. *Cold Spring Harbor Symp. Quant. Biol.* 55:131–136.

Brown, T.H., E.W. Kairiss, and C.L. Keenan. 1990. Hebbian synapses: Biophysical mechanisms and algorithms. *Ann. Rev. Neurosci.* 13:475–511.

Burke, R.E. 1990. Spinal cord: ventral horn. In: Synaptic Organization of the Brain, G. Shepherd, ed., pp. 88–132, 3rd edition. Oxford: Oxford Univ. Press.

Burnod, Y., and H.Korn. 1989. Consequences of stochastic release of neurotransmitters for network computation in the central nervous system. *Proc. Natl. Acad. Sci.* 86:352–356.

---

[2] In other words, to what extent is I_A in hippocampal pyramidal cells necessary to understand memory storage and retrieval? Maybe yes, maybe no—we do not really know.

Bush, P., and R.J. Douglas. 1991. Synchronization of bursting action potential discharge in a model network of cortical neurons. *Neural Comp.* 3:19–30.

DeWeerth, S.P., and C. Mead. 1990. An analog VLSI model of adaptation in the vestibulo-ocular reflex. In: Neural Inf. Pro. Systems 2, D. Touretzky, ed., pp. 742–749. San Mateo, CA:Kaufmann.

Douglas, R., and K. Martin. 1991. Opening the grey box. *TINS* 14:286–293.

Egelhaaf, M., K. Hausen, W. Reichardt, and C. Wehrhahn. 1988. Visual course control in flies relies on neuronal computation of object and background motion. *TINS* 11:351–358.

FitzHugh, R. 1961. Impulses and physiological states in theoretical models of nerve membrane. *Biophys. J.* 1:445–466.

Fogelson, A.L., and R.S. Zucker. 1985. Presynaptic calcium diffusion from various arrays of single channels. *Biophys. J.* 48:1003–1017.

Guttman, R., S. Lewis, and J. Rinzel. 1980. Control of repetitive firing in squid axon membrane as a model for a neuron oscillator. *J. Physiol.* 305:377–395.

Hooper, S.L., and Marder, E. 1987. Modulation of a central pattern generator by the peptide proctolin. *J. Neurosci.* 7:2097–2112.

Hopfield, J.J. 1984. Neurons with graded response have collective computational properties like those of two-state neurons. *Proc. Natl. Acad. Sci.* 81:3088–3092.

Jack, J.J.B., A.U. Kullmann, A.U. Larkman, G. Major, and K.J. Stratford. 1990. Quantal analysis of excitatory synaptic mechanisms in the mammalian central nervous system. *Cold Spring Harbor Symp. Quant. Biol.* 55:57–67.

Katz, B. 1966. Nerve, Muscle and Synapse. New York: McGraw-Hill.

Kepler, T.B., E. Marder, and L.F. Abbott. 1990. The effect of electrical coupling on the frequency of model neuronal oscillators. *Science* 248:83–85.

Koch, C., R. Douglas, and U. Wehmeier. 1990. Visibility of synaptically induced conductance changes: theory and simulations of anatomically charecterized cortical pyramidal cells. *J. Neurosci.* 10:1728–1744.

Koch, C., and T. Poggio. 1987. Biophysics of computation: Neurons, synapses and membranes. In: Synaptic Function, ed. G.M. Edelman, W.E. Gall, and W.M. Cowan, pp. 637–697. New York: Wiley.

Koch, C., T. Poggio, and V. Torre. 1982. Retinal ganglion cells: A functional interpretation of dendritic morphology. *Phil. Trans. R Soc. Lond. B.* 298:227–264.

Korn, H., and D.S. Faber. 1987. Regulation and significance of probabilistic release mechanisms at central synapses. In: Synaptic Function, ed. G.M. Edelman, W.E. Gall, and W.M. Cowan, pp. 57–107. New York:Wiley.

Korn, H., and D.S. Faber. 1991. Quantal analysis and synaptic efficacy in the CNS. *TINS* 14:439–445.

Korn, H., A. Triller, A. Mallet, and D.S. Faber. 1981. Fluctuating responses at a central synapse: *n* of binominal fit predicts number of stained presynaptic boutons. *Science* 213:898–901.

Larkman, A., K. Stratford, and J. Jack. 1991. Quantal analysis of excitatory synaptic action and depression in hippocampal slices. *Nature* 350:344–347.

Lazzaro, J., and C.A. Mead. 1989. A silicon model of auditory localization. *Neural Comp.* 1:47–57.

Lehky, S.R., and T.J. Sejnowski. 1988. Network model of shape-from-shading: Neural function arises from both receptive and projective fields. *Nature* 333:452–454.

Llinas, R. 1988. The intrinsic electrophysiological properties of mammalian neurons: Insights into central nervous system function. *Science* 242:1654–1664.

Llinas, R., and M. Sugimore. 1980. Electrophysiological properties of *in vitro* Purkinje cell dendrites in mammalian cerebellar slices. *J. Physiol.* 305:197–213.

Lustig, G., H. Parnas, and L.A. Segel. 1989. Neurotransmitter release: development of a theory for total release based on kinetics. *J. Theor. Biol.* **136**:151–170.

Lytton, W.W., and T. Sejnowski. 1991. Simulations of cortical pyramidal neurons synchronized by inhibitory interneurons. *J. Neurophysiol.*, in press .

Mahowald, M., and R. Douglas. 1991. The silicon neuron. *Nature*, in press .

Marder, E., L.F. Abbott, F. Buchholtz, I.R. Epstein, J. Golowasch, S.L. Hooper, and T.B. Kepler. 1992. Physiological insights from cellular and network models of the stomatogastric nervous systems of lobsters and crabs. *Am. Zool.*, in press.

Marr, D., and T. Poggio. 1977. From understanding computation to understanding neural circuitry. *Neurosci. Res. Prog. Bull.* **15**:470–488.

McClelland, J.L., and D.E. Rumelhart. 1988. Explorations in Parallel Distributed Processing. Cambridge, MA: MIT Press.

McCormick, D. 1990. Membrane properties and neurotransmitter actions. In: The Synaptic Organization of the Brain, ed. G. Shepherd, pp. 220–243. Oxford: Oxford Univ. Press.

McCormick, D., and H.C. Pape. 1990. Noradrenergic and serotonergic modulation of a hyperpolarization-activated cation current in thalamic relay neurons. *J. Physiol.* **431**:319–342.

Mead, C. 1989. Analog VLSI and neural systems. Reading, MA: Addison-Wesley.

Mead, C., and L. Conway. 1980. Introduction to VLSI Systems. Reading, MA: Addison-Wesley.

Mead, C., and M.A. Mahowald. 1988. A silicon model of early visual processing. *Neural Networks* **1**:91–97.

Mel, B. 1992. NMDA-based pattern discrimination in a modeled cortical neuron. *Neural Comp.*, in press.

Miles, R., and R.K.S. Wong. 1983. Single neurones can initiate synchronized population discharge in the hippocampus. *Nature* **306**:371–373.

Miles, R., and R.K.S. Wong. 1986. Excitatory synaptic interactions between CA3 neurons in the guinea pig hippocampus. *J. Physiol.* **373**:397–418.

Nagumo, J.S., S. Arimoto, and S. Yoshizawa. 1962. An active pulse transmission line simulating nerve axon. *Proc. IRE* **50**:2061–2070.

Pape, H.C., and D.A. McCormick. 1989. Noradrenaline and serotonin selectively modulate thalamic burst firing by enhancing a hyperpolarization-activated cation current. *Nature* **340**:715–718.

Parnas, H., J. Dudel, and I. Parnas. 1986. Neurotransmitter release and its facilitation in crayfish. VII. Another voltage dependent process besides Ca^{2+} entry controls the time course of neurotransmitter release. *Pflügers Arch.* **406**:121–130.

Rall, W. 1989. Cable theory for dendritic neurons. In: Methods in Neuronal Modeling, ed. C. Koch and I. Segev, pp. 9–62. Cambridge, MA: MIT Press.

Rall, W., and G.M. Shepherd. 1968. Theoretical reconstruction of field potentials and dendrodendritic synaptic interactions in olfactory bulb. *J. Neurophysiol.* **31**:884–915.

Rall, W., G.M. Shepherd, T.S. Reese, and M.W. Brightman. 1966. Dendrodendritic synaptic pathway for inhibition in the olfactory bulb. *Exp. Neurol.* **14**:44–56.

Redman. S.J. 1990. Quantal analysis of synaptic potentials in neurons of the central nervous system. *Physiol. Rev.* **70**:165–198.

Reichardt, W., T. Poggio, and K. Hausen. 1983. Figure-ground discrimination by relative motion in the visual system of the fly. Part II: towards the neural circuitry. *Biol. Cybern.* **46** (Suppl.):1–30.

Rinzel, J., and G.B. Ermentrout. 1989. Analysis of neural excitability and oscillations. In: Methods in Neuronal Modeling, ed. C. Koch and I. Segev, pp. 135–169. Cambridge, MA: MIT Press.

Ryckebusch, S., J.M. Bower, and C. Mead. 1989. Modeling small oscillating biological networks in analog VLSI. In: Neural Inf. Proc. Systems 1, ed. D.S. Touretzky, pp. 384–393. San Mateo. CA: Kaufmann.

Sejnowski, T., C. Koch, and P.S. Churchland. 1988. Computational neuroscience. *Science* **241**:1299–1306.

Sharp, A.A., L.F. Abbott, and E. Marder. 1992. Artificial electrical synapses in oscillatory networks, submitted for publication.

Shepherd, G.M. 1972. The neuron doctrine: A revision of functional concepts. *Yale J. Biol. Med.* **45**:584–599.

Simon, S.M., and R.R. Llinas. 1985. Compartmentalization of the submembrane calcium activity during calcium influx and its significance in transmitter release. *Biophys. J.* **48**:485–498.

Stratford, K., A. Mason, A. Larkman, G. Major, and J. Jack. 1989. The modeling of pyramidal neurones in the visual cortex. In: The Computing Neuron, ed. R. Durbin, C. Miall, and G. Mitchison, pp. 296–321. Reading, MA: Addison-Wesley.

Traub, R., and R. Miles. 1991. Neuronal Networks of the Hippocampus. Cambridge, MA: Cambridge Univ. Press.

Traub, R.D., and R.K.S. Wong. 1982. Cellular mechanism of neuronal synchronization in epilepsy. *Science* **216**:745–747.

Traub, R.D., R.K.S. Wong, R. Miles, and H. Michelson. 1991. A model of CA3 hippocampal pyramidal neuron incorporating voltage-clamp data on intrinsic conductances. *J. Neurophysiol.* **66**:635–650.

Wilson, M.A., and J.M. Bower. 1989. The simulation of large-scale neural networks. In: Methods in Neuronal Modeling, ed. C. Koch and I. Segev, pp. 291–334, Cambridge, MA: MIT Press.

Wörgötter, F., and C. Koch. 1991. A detailed model of the primary visual pathway in the cat: Comparison of afferent excitatory and intracortical inhibitory connection schemes for orientation selectivity. *J. Neurosci.* **11**:1959–1979.

Yamada, W.M., D. Koch, and P.R. Adams. 1989. Multiple channels and calcium dynamics. In: Methods in Neuronal Modeling, ed. C. Koch and D. Seger, pp. 97–133. Cambridge, MA: MIT Press.

Zipser, D., and R.A. Andersen. 1988. A back-propagation programmed network that simulates response properties of a subset of posterior parietal neurons. *Nature* **331**:679–684.