Joint Conference on Neural Networks, vol. 2, New York: IEEE, pp. 357–363.

Rumelhart, D. E., Hinton, G. E., and Williams, R. J., 1986, Learning internal representations by error propagation, in *Parallel Distributed Processing: Explorations in the Microstructure of Cognition*, vol. 1, *Foundations*, (D. E. Rumelhart, J. L. McClelland, and PDP Research Group, Eds.), Cambridge, MA: MIT Press, chap. 8. ◆

Rosenblatt, F., 1962, *Principles of Neurodynamics: Perceptrons and the Theory of Brain Mechanisms*, Washington, DC: Spartan.

Widrow, B., and Hoff, M. E., Jr., 1960, Adaptive switching circuits, in *1960 IRE WESCON Convention Record*, Part 4, New York: IRE, pp. 96–104.

Widrow, B., and Lehr, M. A., 1990, 30 years of adaptive neural networks: Perceptron, madaline, and backpropagation, *Proc. IEEE*, 78:1415–1442. ◆

Widrow, B., Rumelhart, D., and Lehr, M. A., 1994, Neural networks: Applications in industry, business, and science, *Commun. ACM*, 37(3):93–105.

Widrow, B., and Stearns, S. D., 1985, *Adaptive Signal Processing*, Englewood Cliffs, NJ: Prentice-Hall. ◆

Perspective on Neuron Model Complexity

Wilfrid Rall

Introduction

There is a wide range of choice in model complexity, from very simple to rather complex neuron models. Which model to choose depends, in each case, on the context. How much information do we already have about the neurons under consideration? What questions do we wish to explore?

Sometimes we wish to model a particular biological neuron whose anatomy and physiology are known in considerable experimental detail. In such cases, we may choose to specify a model that includes at least some of the dendritic branching of the neuron, because synapses from one source may be distributed preferentially to either a distal or a proximal dendritic location, while synapses from another source may end mainly at the soma, or on a different dendritic tree of the same neuron. Also, there may be a functionally significant nonuniformity in the distribution of channel densities of several ion channel types over the surface of the soma and dendrites. How much detail is needed depends on the biological experiments to be simulated and the questions asked.

Conversely, many network modelers are not constrained by anatomical or physiological data. For some network modeling, this is partly justified by a paucity of available data. However, more often, network modelers are constrained by their mathematical methods, which lead them to focus on abstract networks composed of extremely simple units. The simplest units are two-state, binary units, analogous to atomic spin, previously studied for condensed matter physics (see, e.g., OPTIMIZATION, NEURAL). Such binary units do not resemble neurons, but they do have a strong appeal for nerve-net modelers, who have generated an extensive literature. That literature lies outside the scope of the current article.

When simple binary units are compared with a dendritic neuron model (especially with nonuniform distributions of synapses and ion channels), it becomes apparent that one dendritic model neuron can perform tasks that would require a network of many simple units to duplicate. For the purpose of machine design, it may seem quite appropriate to consider the trade-offs in cost and flexibility (between the one realistic model and the many binary units), but for functional insights and understanding of biological nervous systems, I freely state my bias for the more realistic neuron models. I do not choose the most complex, in the sense of including all known anatomical and physiological details; I favor an intermediate level of complexity that preserves the most significant distinctions between regions (soma, proximal dendritic, distal dendritic, different trees), especially when further justified by nonuniform distributions of synapses and ion channels (see also Segev, 1992).

The claim is sometimes made that network properties depend primarily on the connectivity between the units, and not on the properties of the units. Although this may be true for some gross network properties, I do not believe it to be true for many of the actual biological networks that perform important, complicated tasks. I regard it as a worthwhile challenge for like-minded neural modelers to provide interesting demonstrations in support of this belief. The challenge is to demonstrate a useful computation or discrimination that can be accomplished with a dendritic neuron model, or a network composed of such models, and then show that this useful capacity is lost when all of the dendritic membrane is lumped with the soma, and all of the inputs to each neuron are now delivered to that lumped membrane. There are valuable examples that already meet this challenge, several of which are presented in three later sections of this article. Other examples can be found in a review by Borst and Egelhaaf (1994; see also VISUAL COURSE CONTROL IN FLIES).

Brief Historical Notes

Neurons are biological cells, and their electrical properties depend on ions and the cell membrane, in a manner brilliantly elucidated by Hodgkin, Huxley, and Katz during the period 1948–1952. It is a fascinating historical coincidence that two seeds of their important insights can be found in a single 1902 volume of *Pfluegers Archiv*, in pioneering articles by Bernstein and by Overton. Following the earlier theoretical insights of Nernst and Planck, Bernstein recognized the importance of the potassium ion concentration difference across the membrane in determining a non-zero resting potential; he regarded excitation as a brief breakdown of the membrane, a concept that prevailed until 1948, when Hodgkin and Katz showed that the key is a sudden overwhelming increase in membrane permeability to sodium ions. Overton's 1902 paper had correctly emphasized the importance of the external sodium ion concentration to the excitability properties of nerve, but no one put these ideas together in 1902. Between 1900 and 1914, several investigators, including Hermann, Lucas, and Lapique, recognized the importance of membrane capacitance; the concept of nerve membrane as a leaky integrator, with a threshold for an action potential, was used to understand the strength-duration curve for a threshold stimulus. During the 1930s, several investigators, including Rashevsky, Hill, and Monnier, developed mathematical models of excitation and inhibition; Rashevsky's textbook *Mathematical Biophysics* (1948) includes many examples of network modeling by himself; by Householder, Landahl, and others; and by McCulloch and Pitts, whose famous 1943 paper arose in the context of Rashevsky's research seminars at the University of Chicago (see also the historical notes in Schwartz, 1990). Ever since that time, many neuron modelers have been content with the leaky integrator neuron model, which reduces a neuron to a single node that

integrates synaptic excitation (+) and synaptic inhibition (−) delivered to it by other neurons. Several errors caused by these oversimple assumptions were demonstrated by compartmental computations in 1962; see Rall's chapter in Reiss (1964) or in Segev, Rinzel, and Shepherd (1995). Other chapters in Reiss (1964) also provide several interesting early perspectives on neural modeling. The mathematical modeling of nonlinear membrane properties has been presented and discussed in an outstanding early review by FitzHugh (1969), and in a chapter by Rinzel and Ermentrout that appears in Koch and Segev (1989).

The concept of a nerve axon as an extended core conductor (i.e., membrane cylinder with ionic conducting media inside and outside) goes back to the 1870s, when it was treated mathematically by Hermann and Weber; both the concept of passive electrotonus in membrane cylinders and the mathematics (of passive cable theory) were explored over the years, culminating in classic papers by Hodgkin and Rushton and by Davis and Lorente de Nó, both around 1946–1947; see references in Rall (1977). Before 1900, neuroanatomical studies by Ramón y Cajal demonstrated the extensiveness of dendritic branching for most neuron types; this was confirmed by many anatomists, and later (in the 1950s), use of the electron microscope made it possible to verify the existence of very many synapses on the dendritic branches and on the dendritic spines of neurons. These anatomical facts, together with the introduction of intracellular microelectrode recording from single dendritic neurons (in the 1950s), made it urgent to extend cable theory to the dendrites of individual neurons. This was begun in the late 1950s and carried forward into the 1960s and 1970s; for a review, see Jack, Noble, and Tsien (1975) or Rall (1977); see also Koch and Segev (1989), McKenna, Davis, and Zornetzer (1992), Rall et al. (1992), Segev et al. (1995), and DENDRITIC PROCESSING.

Dendritic Neuron Model Complexity: Geometric Versus Membrane Complexity

The concept of complexity in dendritic neuron models can be explored quite efficiently by making a two-dimensional chart. One dimension would be membrane complexity, ranging from the simple case of a passive linear membrane to that of postsynaptic membrane models with time-varying ion permeability (or conductance), and then to excitable membrane models with voltage-dependent ion conductances as described by Hodgkin and Huxley (see AXONAL MODELING), or as now described with increasing detail in terms of many different species of ion channels whose voltage and time dependence are currently being characterized (see ION CHANNELS: KEYS TO NEURONAL SPECIALIZATION). The other dimension would be geometric complexity, ranging from the simple case of an isopotential region of membrane (a soma, or a space-clamped section of a cylinder) to that of a uniform membrane cylinder with two sealed ends (or with one end voltage clamped, or current clamped), and then to several dendritic trees attached to a soma (with or without an axon), where the soma may be shunted and the branching of the trees may be specified to varying degrees of arbitrariness. The most complicated geometric case, with arbitrary branching and shunted soma, has recently been solved analytically (for transients, assuming passive membrane) in a mathematical tour de force by Major, Evans, and Jack (1993); see also Holmes, Segev, and Rall (1992). The less complicated, but illuminating, case of idealized branching with a point soma was solved earlier by Rall and Rinzel; see the 1973 and 1974 papers reprinted in Segev et al. (1995). However, these analytical methods do depend on the assumption of linear membrane properties. When nonlinear membrane complexity is combined with geometric complexity, the transient solutions can be obtained computationally by using compartmental models; see 1964 and 1968 papers reprinted in Segev et al. (1995); see also DENDRITIC PROCESSING and several chapters in Koch and Segev (1989) and in McKenna et al. (1992).

Dendritic Model Can Provide Spatiotemporal Discrimination

Figure 1 summarizes a demonstration of how a dendritic neuron model could perform a discrimination between two contrasting spatiotemporal patterns of synaptic input (i.e., possible movement detection); this discrimination is lost if the compartments and inputs are lumped together. A neuron is represented by a chain of ten compartments; compartment 1 represents the soma, while compartments 2 to 10 represent dendritic membrane of the same neuron, with increasing cable distance from the soma. One spatiotemporal input sequence, A-B-C-D, has the proximal dendritic input first, followed in time by progressively more distal dendritic input locations. The other input pattern, D-C-B-A, is opposite in having the most distal input first, followed in time by progressively more proximal input locations. Comparison of the resulting computed voltage transients (EPSP at the soma), shown in Figure 1, reveals that input sequence D-C-B-A yields a significantly larger voltage amplitude than does input sequence A-B-C-D. Intuitive understanding of this computed result is obtained by noting that the delayed proximal input builds on membrane depolarization that has spread to the soma (with delay) from the distal dendrites (which were activated earlier). If the voltage threshold for spiking at the soma were tuned between these two peak amplitudes, a spike would be produced by sequence D-C-B-A, but a spike would not be produced by sequence A-B-C-D; this would constitute a discrimination between these two sequences. The dashed curve in the figure shows the computed result when the compartments are lumped together; either sequence of input synapse activation then produces the same intermediate result, and no discrimination would be possible.

Models for Mitral and Granule Cell Populations in Olfactory Bulb

A rather different example is provided by the neuron models used for the mitral cell and granule cell populations in simulating experiments on the OLFACTORY BULB (q.v.) of rabbit; see the 1968 paper of Rall and Shepherd in Segev et al. (1995); or see figures 2.11 and 2.12 in Koch and Segev (1989). Here, the task was to model and compute extracellular field potentials that matched those observed experimentally in olfactory bulb when the mitral cell population was activated in near synchrony by means of an antidromic volley. Compartmental models were used; a nine-compartment model (three axonal, one somatic, and five dendritic) was used to simulate antidromic activation of a mitral cell, while a ten-compartment model was used to simulate nonspiking activity in the dendrites of an axonless granule cell. The dendritic compartments were absolutely essential for the computation of electric current flow between different dendritic regions of each granule cell and between the dendrites and soma of each mitral cell; without these currents, it would have been impossible to compute the field potentials generated by the synchronously activated neuron populations. Also, this modeling led to a critically important distinction in the depth distribution of the two fields: the larger, longer-lasting field potentials generated by the very large population of granule cells extended to significantly greater depth in the olfactory bulb than did the earlier, smaller, briefer field potentials generated by the mitral cell population. The difference between these two fields was such that neither population could have generated the other field. This provided the key to our prediction of (and the functional interpretation of subsequent electron microscopic evidence for) dendrodendritic synaptic interactions between the mitral secondary dendrites and the distal dendrites of the granule cells, which are

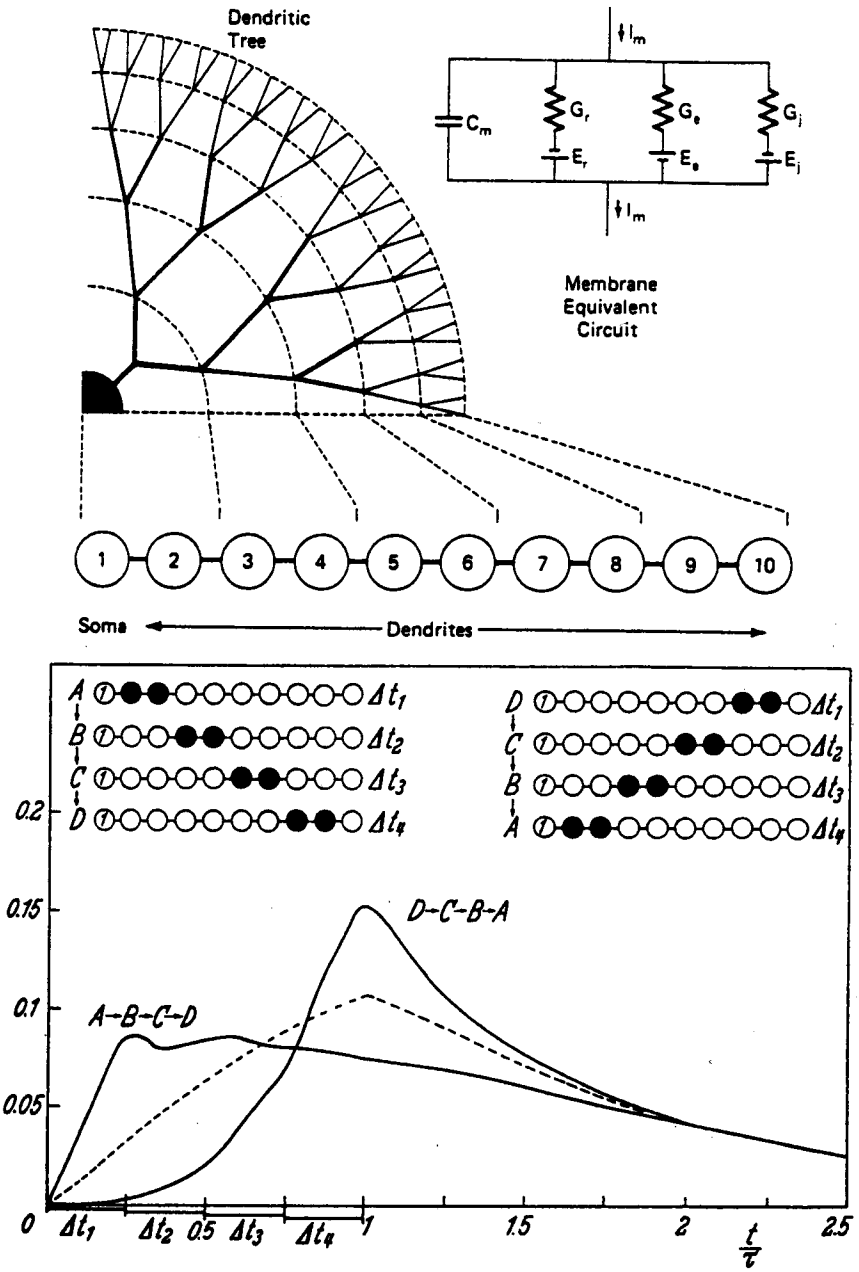

Figure 1. Effect of spatiotemporal dendritic pattern of synaptic input on the computed EPSP at the soma, for a ten-compartment model. Upper diagram indicates the mapping of a soma and dendritic tree into a chain of ten equal compartments. Compartment 1 represents the soma membrane, while compartments 2 to 10 represent dendritic membrane, from proximal to distal locations. The middle diagram (at left) shows the synaptic input sequence A-B-C-D, meaning proximal dendritic input location active first, followed by successive activation at increasingly more distal input locations; this input pattern produced the soma voltage transient (computed composite EPSP) labeled A-B-C-D at lower left. The middle diagram (at right) shows the opposite synaptic input sequence. D-C-B-A, meaning distal dendritic input location first, followed by successively more proximal input locations; this input pattern produced a significantly different soma voltage transient (computed composite EPSP), having a delayed rise to a larger peak amplitude, labeled D-C-B-A. In both cases, each input compartment (shown in black) received a synaptic excitatory conductance pulse ($G_e = G_r$, for a duration 0.25τ) during one of the four labeled periods. The same total amount of synaptic input produced the dashed curve when the spatiotemporal pattern was eliminated by smearing the synaptic conductance in space and time ($G_e = 0.25\ G_r$ in eight compartments (compartments 2 to 9) for the full time duration from $t = 0$ to $t = \tau$). The membrane equivalent circuit (upper right) holds for each compartment. Further details can be found in the 1964 chapter by Rall in Reiss (1964), reprinted in Segev et al. (1995).

intermingled in the external plexiform layer of the bulb. If these cells had been modeled as lumped somas, without dendrites, neither the successful simulation of the experimental field potentials nor the exciting new insights about a dendrodendritic pathway for recurrent inhibition would have been possible.

Similarly, for the earlier simulations and insights obtained for motor neurons of cat spinal cord, we found that observations made at the soma seemed paradoxical until they were understood in terms of synaptic events that occur in distal dendrites (see the 1967 paper in Segev et al., 1995); these results and insights would not have been possible without dendritic compartments in the neuron field.

Comment on Functional Aspect of Dendrodendritic Interactions

To highlight an important functional difference, note first that motor neurons do exhibit the classical functional polarity envisaged by Ramón y Cajal and Sherrington (as well as most modelers). The dendrites receive inputs from many sources (their effects converge on the soma); the output (when spike threshold is exceeded) is an all-or-nothing action potential propagated by the axon to muscle units that may be quite distant; i.e., classically, input is received by the dendrites and output is delivered by the axon. In contrast, the dendrites of both the mitral cells and the granule cells are functionally different, because they both send as well as receive synaptic information, locally. The mitral secondary dendrites, which are smooth and spineless, send nonspiking synaptic excitatory output, which is received as input by the spines (see DENDRITIC SPINES) of the adjacent granule cells. The granule cells have no axons and perhaps no action potentials; their spines receive graded synaptic excitatory input and then send graded synaptic output that is inhibitory to the adjacent mitral cell dendrites. It is important to emphasize that this is not a rare anomaly found only in the olfactory

bulb; evidence for dendrodendritic synapses and for graded local synaptic interactions is now found in many parts of the brain (e.g., retina and inferior olive). In 1965, when we (Rall et al.; see 1966 and 1968 papers reprinted in Segev et al., 1995) first presented our interpretations of dendrites that send as well as receive, some critics resisted this concept as heretical; however, our functional interpretation of these dendrodendritic synapses is now widely accepted by physiologists and anatomists. This kind of graded two-way synaptic interaction is very different from the classical functional polarity just described for motor neurons; it provides graded functional coupling between neurons (without axonal impulses). The implications have so far hardly been explored in theoretical networks. Such exploration will require explicit modeling of dendritic compartments; a point neuron model would be useless for this. Note also that computational exploration of localized plastic changes at synapses and at dendritic spines depends on neuron models that include dendritic compartments.

Network Rhythmogenesis Using the Traub Model and a Reduced Model

A 19-compartment cable model for the pyramidal cells of the CA3 region of guinea pig hippocampus was developed by Traub et al. (1991; see also the chapter by Traub and Miles in McKenna et al., 1992). Based on experimental measurements, parameters were chosen for each compartment, using up to six active ionic conductances, and controlled by ten-channel gating variables. They succeeded in finding a set of physiologically reasonable parameters for which the network of model neurons could simulate several important aspects of the experimental repertoire of the slightly disinhibited hippocampal slice preparation. Although Traub et al. recognized that their successful simulations of network behavior depended on specifying significantly different ion channel densities for the soma and for the dendrites, the critical importance of this difference was made starkly clear by the modeling of Pinsky and Rinzel (1994); they obtained essentially the same behavioral repertoire by using a network composed of a severely reduced neuron model consisting of only two compartments per pyramidal cell. One compartment represented the soma and proximal dendrites, while the other compartment represented the distal dendrites. To be more specific, the ion channels for fast-spiking currents (inward sodium, and delayed rectifier) were restricted to the soma-like compartment, and the ionic channels for the slower calcium currents (calcium-inward and calcium-modulated currents) were restricted to the dendrite-like compartment. I hasten to add that these results also show that at least two compartments are needed for simulations of this behavior; a single lumped compartment, with all of the ion channels in parallel, could not produce the same behavior, especially the rhythm, which basically involves an alternating flow of current between the two coupled compartments. A special advantage of the reduced neuron model is that much simpler computations can explore how much the interesting behavior depends on the values of key parameters, especially the parameter that defines the tightness of coupling between the two compartments. Also, the behavior of very large networks can be explored more efficiently using such a reduced neuron model. Further study may show that the two-compartment model cannot match the fuller model in certain important tests, but, in any case, these findings so far represent a very satisfying example that illustrates the thesis of this article.

Discussion

In an earlier essay offering perspective on neural modeling (a chapter in Binder and Mendell, 1990), I provided a completely different set of examples. One of these provided a detailed consideration of the number of degrees of freedom to be found in a neuron model composed of a thousand compartments. Such models exist today because of tremendous improvements in anatomical methods and in computation facilities now available to experimental investigators. Because they have the morphological data and a computer, why not put everything into the model? The answer is that you can if you wish to, but you should be aware of the huge number of degrees of freedom implied by the large number of parameters that must be specified; as someone once pointed out, given enough free parameters, he could fit an elephant. Is the membrane uniform, or do we know the density of every channel species in every membrane compartment? How are the inputs distributed to the many compartments? Today, the data needed for such detailed specifications are largely missing; however, such data are beginning to become at least partly available for some neurons. Where the data are not available, the modeler must make reasonable guesses. If it seems reasonable to assign the same parameter values to many neighboring compartments, one should consider lumping those compartments together to produce a simpler model with fewer compartments. Nevertheless, one important merit of the larger model is that it can be used to test whether it can perform some interesting task that cannot be performed by the reduced model.

As stated earlier, my preference is for intermediate levels of complexity; I vote for the smallest number of compartments that can preserve what one judges to be the functionally important differences between dendritic regions with regard to ion channel densities and to distributions of synapses from different sources. If a five-compartment model can provide a good approximation of the interesting properties of a 1,000-compartment model, I would prefer the smaller model, for two important reasons: (1) it helps sharpen our intuitive understanding about what is essential to obtaining the behavior of interest, and (2) it can greatly facilitate computations with networks composed of such neuron models. I expect modeling of this kind will continue to be particularly fruitful in the near future (see also the discussion by Segev, 1992).

Concluding Comment

As when drawing, painting, sculpting, or composing music, so too, when deeply engaged in neural modeling, I believe that much of the fun and satisfaction comes from interactions between my conscious mind and my subconscious sources of creativity. It seems that preliminary sketching serves to plant seeds in the subconscious, where they can grow, if nurtured. Conscious pursuit of the problem can then stimulate differentiation and development in the subconscious and may produce fruits that can reach conscious awareness (popping up like mushrooms produced by an underground mycelium). Such fruits may provide exciting new insights for the conscious mind. Indeed, the pleasure of such creative discovery can become almost addictive for those fortunate enough to have both the interest and the opportunity for creative activity. I hasten to add that a lot of hard work is usually required to test and polish before one can produce a finished product. Pioneering in dendritic neuron modeling provided me with such an opportunity; now [at the time of the First Edition], with retirement upon me, I hope to persist by sculpting, painting, and by designing a house for a natural mountain setting.

[Reprinted from the First Edition]

Road Maps: Biological Neurons and Synapses; Grounding Models of Neurons
Background: I.1. Introducing the Neuron
Related Reading: Dendritic Processing

References

Binder, M. D., and Mendell, L. M., Eds., 1990, *The Segmental Motor System*, Oxford: Oxford University Press.

Borst, A., and Egelhaaf, M., 1994, Dendritic processing of synaptic information by sensory interneurons, *Trends Neurosci.*, 17:257–263.
FitzHugh, R., 1969, Mathematical models of excitation and propagation in nerve, in *Biological Engineering* (H. P. Schwann, Ed.), New York: McGraw-Hill. ◆
Holmes, W. R., Segev, I., and Rall, W., 1992, Interpretation of time constant and electrotonic length estimates in multi-cylinder or branched neuronal structures, *J. Neurophysiol.*, 68:1401–1420.
Jack, J. J. B., Noble, D., and Tsien, R. W., 1975, *Electric Current Flow in Excitable Cells*, Oxford: Oxford University Press. ◆
Koch, C., and Segev, I., 1989, *Methods in Neuronal Modeling: From Synapses to Networks*, Cambridge, MA: MIT Press. ◆
Major, G., Evans, J. D., and Jack, J. J. B., 1993, Solutions for transients in arbitrarily branching cables: I. Voltage recording with a somatic shunt, *Biophys. J.*, 65:423–449.
McKenna, T., Davis, J., and Zornetzer, S. F., 1992, *Single Neuron Computation*, San Diego, CA: Academic Press. ◆
Pinsky, P. F., and Rinzel, J., 1994, Intrinsic and network rhythmogenesis in a reduced Traub model for CA3 neurons, *J. Computat. Neurosci.*, 1:39–60.
Rall, W., 1977, Core conductor theory and cable properties of neurons, in *Handbook of Physiology: The Nervous System: Cellular Biology of Neurons*, sect. 1, vol. I, part 1, chap. 3, Bethesda, MD: American Physiological Society, pp. 39–97. ◆
Rall, W., Burke, R. E., Holmes, W. R., Jack, J. J. B., Redman, S. J., and Segev, I., 1992, Matching dendritic neuron models to experimental data, *Physiol. Rev.*, 72:S159–S186. ◆
Rashevsky, N., 1948, *Mathematical Biophysics*, Chicago: University of Chicago Press; reissued 1960, New York: Dover.
Reiss, R., Ed., 1964, *Neural Theory and Modeling*, Stanford, CA: Stanford University Press.
Schwartz, E. L., 1990, *Computational Neuroscience*, Cambridge, MA: MIT Press. ◆
Segev, I., 1992, Single neurone models: Oversimple, complex and reduced, *Trends Neurosci.*, 15:414–421. ◆
Segev, I., Rinzel, J., and Shepherd, G. M., Eds., 1995, *The Theoretical Foundation of Dendritic Function: Selected Papers of Wilfrid Rall with Commentaries*, Cambridge, MA: MIT Press. ◆
Traub, R., Wong, R., Miles, R., and Michelson, H., 1991, A model of a CA3 hippocampal pyramidal neuron incorporating voltage-clamp data on intrinsic conductances, *J. Neurophysiol.*, 66:635–649.

Phase-Plane Analysis of Neural Nets

Bard Ermentrout

Introduction

Models of neural networks often involve the solutions to differential equations that describe the time evolution of these complex systems. The dynamical behavior of these networks ranges from the convergence to an equilibrium (generally desired in connectionist applications) to oscillatory behavior (in models of central pattern generators and bursting) through possibly chaotic behavior. There are many ways to analyze these models; the most commonly used techniques entail simulation. In this article I will give an overview of an alternative technique for studying the *qualitative* behavior of small systems of interacting neural networks. One form that the models take is (Ellias and Grossberg, 1975; Hopfield, 1984; Wilson and Cowan, 1972):

$$\tau_i \frac{dx_i}{dt} = -x_i + f_i\left(\sum_{j=1}^{n} w_{ij}x_j + s_i\right) \quad i = 1, \ldots, n \qquad (1)$$

where x_i represents the activity or firing rate of the ith neuron, τ_i is the time constant, w_{ij} are the connection weights, s_i are inputs, and f_i are typically saturating nonlinear functions that have the form shown in Figure 1. That is, the nonlinear functions are increasing and bounded. Some typical examples are:

$$f(x) = \tanh(x) \qquad (2)$$

$$f(x) = \tan^{-1}(x) \qquad (3)$$

$$f(x) = \frac{1}{1 + \exp(-x)} \qquad (4)$$

Often, a slightly different form of (1) is chosen where the nonlinearities are inside the sums. The transformation from one to the other is elementary and all of the following holds for either type of model.

A complete analysis of networks of the form in Equation 1 is obviously impossible. However, if $n \leq 2$, then a fairly complete description of Equation 1 can be given. Thus, the goal of this article is to introduce the reader to the qualitative theory of differential equations in the plane. In particular, I will analyze two neuron networks that consist of (1) two excitatory cells, (2) two inhibitory cells, and (3) an excitatory and an inhibitory cell. The advantages of restricting the analysis to these small networks are the special topology of the plane, the completeness of the analysis possible, and finally the case of exposition. Indeed, an overview of nonlinear dynamics can be obtained through these simple examples. Beer (1995) has attempted to exhaustively study the dynamics in the case $n = 2$ and gives a nearly complete overview of the possible types of behavior that can be expected. However, he does miss several interesting examples (Ermentrout, 1998, pp. 371–373). Another more general approach for the analysis of large numbers of coupled systems is to use bifurcation methods that enable one to *reduce* the dimensionality of the resulting equations and then apply techniques such as those used here. While planar systems may seem to be a rather extreme simplification, there is some justification for it. For example, in some local cortical circuits, there is no structure in the connectivity and there are essentially two types of neurons, excitatory and inhibitory. Thus, we can view the simple planar system as representing a population of coupled excitatory and inhibitory neurons. This approach was used successfully to study cortical processing in the rodent somatosensory system (Pinto et al., 1996) and to explain the effects of altering inhibitory interneurons in the hippocampus (Tsodyks et al., 1997).

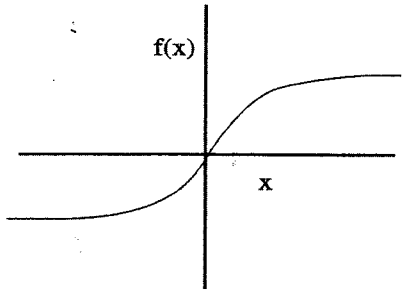

Figure 1. Typical nonlinear input-output function of a single model neuron.