Information Processing in the Dendritic Net

Bruce MacLennan[†]

CS-92-180

Abstract

The goal of this paper is a model of the dendritic net that: (1) is mathematically tractable, (2) is reasonably true to the biology, and (3) illuminates information processing in the neuropil. First I discuss some general principles of mathematical modeling in a biological context that are relevant to the use of linearity and orthogonality in our models. Next I discuss the hypothesis that the dendritic net can be viewed as a linear field computer. Then I discuss the approximations involved in analyzing it as a dynamic, lumped-parameter, linear system. Within this basically linear framework I then present: (1) the self-organization of matched filters and of associative memories; (2) the dendritic computation of Gabor and other nonorthogonal representations; and (3) the possible effects of reverse current flow in neurons.

---

*Based on a presentation at the 2nd Annual Behavioral and Computational Neuroscience Workshop, Georgetown University, Washington DC, May 18–20, 1992.

[†]Computer Science Department, University of Tennessee, Knoxville TN 37996; maclennan@cs.utk.edu.

1 Goal

My goal is a model of the dendritic net that:

- is mathematically tractable,
- is reasonably true to the biology, and
- illuminates information processing in the neuropil.

The approach is different from previous approaches, which lack at in least one of these desiderata. For example, at least from the time of Hodgkin and Huxley we have had differential equations that describe the dynamics of ion flows across neural membranes. Unfortunately, these equations are much too complex and at much too low a level to tell us much about the representation and processing of information in the brain. At the other extreme, we have the neural network models, beginning with McCulloch and Pitts and continuing through modern PDP models. Although these have a relatively tractable mathematical theory and are at a high enough level to explicate information processing, they deviate from the biology in many important respects (Crick & Asanuma 1986; Shepherd 1990a). First, by treating dendrites as no more than the input wires to neurons, they ignore the important role that the dendrites play in neural information processing. Second, by treating neurons as simple discrete time systems with no spatial structure, they ignore much of the important spatiotemporal signal processing that occurs in the nervous system.[1] In this I follow Shepherd (1990, p. 20):

> The vast majority of neural network simulations — in particular, connectionist models — consider individual nerve cells to be single node, linear integration devices. They thus neglect the effect of dendritic, synaptic, and intrinsic membrane properties on the function of individual cells. An important goal of the study of synaptic organization is therefore to identify the specific operations that arise from these properties and incorporate them into more realistic network simulations of specific brain regions.

[1] See Roberts & Bush (1981), especially the papers by Bullock, Rall and Shepherd.

2 Some Principles of Mathematical Modeling

Before describing the model, I suggest two principles to guide the mathematical modeling of the nervous system, the *Complementarity Principle* and the *Robustness Principle*. Like other scientific principles, such as the Verification Principle, these are neither synthetic nor analytic, and so they can be established neither empirically nor deductively. Rather, they should be considered normative principles that must be assessed pragmatically by their long-term consequences for the scientific enterprise.

The Complementarity Principle

> *If your theory predicts different outcomes depending on whether you use discrete mathematics or continuous mathematics, then you've got the wrong theory.*

In simple terms, this principle says that continuous mathematics should be equivalent to discrete mathematics with fine grid, and that discrete mathematics should be equivalent to continuous mathematics with steep transitions. Its practical consequence is that we can move between the two as convenient. Although this may seem obvious enough, it's worthwhile to be reminded of it, since psychology is periodically infected by debates about whether cognition is *really* digital or *really* analog, which are prone to degenerating into medieval scholasticism. The point is, that *mathematical* differences between the discrete and the continuous can't be relevant to our enterprise. What is relevant is whether the phenomena look more discrete or more continuous at our level of analysis. Since all we intend is an approximation to reality, the essential point is that we may use discrete mathematics to approximate the continuous as closely as we like, and vice versa. In stronger terms, we may put this as

The Ontological Argument (or Nobel Prize Argument): If you have a theory that predicts different macroscopic phenomena depending on whether space, time or other physical variables are *really* discrete or *really* continuous, then you have a design for an experiment that will tell us, not just another approximation to the structure of the physical universe, but an absolute and final answer to it. If this is so then you should immediately

abandon whatever you are doing and set up that experiment, because you will almost surely earn a Nobel prize.

In the absence of such a theory the complementarity principle applies: continuous mathematics is equivalent to approximately-continuous, discrete mathematics, and conversely, discrete mathematics is equivalent to approximately-discrete, continuous mathematics. Thus we can move back and forth between the two as convenient.

So much for the Complementarity Principle; consider now:

The Robustness Principle

No biological process can depend on an absolute mathematical property.

This principle may be justified by

The Argument from Noise: Thermal noise and other sources of error and imprecision limit the accuracy of biological processes. For example, if a biological process depended on two things being *exactly* equal or something having *exactly* the correct value, it could never work in a real biological context. Another example: orthogonality, as normally defined in mathematics, is an absolute property: two vectors either are or are not orthogonal; closeness to orthogonality is not relevant.[2] A biological process could not depend on whether two vectors are exactly orthogonal (versus nearly orthogonal). Certainly, given the one or perhaps two digits of precision of axonal signalling, the most we could expect is neuronal processes that depend on orthogonality to within a few percent of their inner products.

One conclusion we can draw from the Robustness Principle is that representations in the brain cannot depend on whether a set of representing functions is an orthogonal basis or a nonorthogonal basis; the most that we can assume is approximate orthogonality. In other words, we're much better off assuming representations are nonorthogonal and then talking about degrees of orthogonality. Another example: the infinite, but vanishingly small, support of the Gabor elementary functions is not a significant problem compared to the finite support of other wavelet families. Finally, and most importantly for the model presented here, we cannot expect a biological system to depend on exact linearity, but approximate linearity is reasonable.

[2] Notions of approximate orthogonality, such as ϵ-orthogonality, have been proposed, most recently by Kainen (1992), and are discussed below.

To emphasize the point of the Robustness Principle, I will digress to consider the precision of computation in the nervous system. First consider signalling by action potential density. The Gabor Uncertainty Principle with 1 kHz bandwidth implies that N msec. are required to distinguish N values (MacLennan 1991). Therefore it takes 10 msec. to transmit an analog value with one digit of precision, and 100 msec. for a value with two digits. Hence we can expect at most 1 to 2 digits precision in axonal signalling.

Next consider the transmission of analog values across a chemical synapse. Since releases of neurotransmitter are quantal, and an action potential releases 100–200 quanta (Shepherd 1988, p. 36), we can expect only about 2 digits precision from a chemical synapse.

Electrical synapses are not well understood and seem to be rare in the mammalian brain. Nevertheless, the quantal change in conductance is about 5 times that of a chemical synapse (120 pS vs. 20–40 pS[3]), suggesting that it is even lower precision (but quicker) (Shepherd, 1988, pp. 70–2, 128–9; Neher, 1992; Sakman, 1992).

Membrane juxtaposition also results in nonsynaptic electrical field potential effects, sometimes called *ephapses*, especially between dendrites (e.g., Shepherd 1988, pp. 123–4, 1990, pp. 48–50). These interactions are difficult to study, but it seems safe to say that since they are very weak (about 10 mV at most), thermal and electrical noise will keep the precision low.

In conclusion I claim:

> *1% to 10% is "close enough for brain work."*

Of course, the foregoing should not be interpreted as an argument against the attempt to find deep mathematical principles underlying neural processes. There are many instances in the history of science where taking the mathematics seriously pointed the way to future discoveries; special relativity and Dirac's postulation of the positron are examples. In the area of neuroscience we can cite the discovery of Gabor receptive fields. On the other hand, there have also been cases where mathematically obvious generalizations have apparently been ignored by Nature (the magnetic monopole?). Thus:

> *We should take our mathematical models seriously —*
> *but not* **too** *seriously.*

[3]The abbreviation pS stands for picosiemen, a unit of conductance equal to 1000 megohm^{-1}.

3 The Dendritic Model of Computation in the Brain

The view I take in this paper is that the neuron is predominantly a transducer, and that computation is implemented by linear processes in the dendritic net. The approach is mainly based on Karl Pribram's (1991) ideas, but also on Gordon Shepherd's notion that the synapse should be viewed as the basic computing element of the brain.

3.1 Shepherd's Microcircuits

Since the mid 1960s Shepherd has stressed the importance of dendritic interactions, and has shown the presence of microcircuits even within a single dendritic spine (e.g., Shepherd 1972, 1978, 1988, 1990). The simplest example is the *reciprocal synapse* that occurs between mitral cells and granule cells in the olfactory bulb. The reciprocal synapse comprises an inhibitory and an excitatory synapse that transmit in opposite directions. Thus a depolarization wave spreading through the mitral-cell dendrite can activate the excitatory synapse to the granule-cell dendrite, thereby causing a depolarization wave to spread though the granule-cell dendrite. The latter depolarization activates the inhibitory synapse, which suppresses the original depolarization of the mitral cell. Thus we have a negative feedback loop implemented within a dendritic spine, about one micron in diameter. The effect is that of an *operational amplifier*, one of the basic components of an analog computer (Fig. 1).

Microcircuits of this kind have been found now in many areas of the brain, including the thalamus, the basal ganglia of the cerebrum and the midbrain, the motor area of the cerebral cortex, the trigeminal nerve, and the suprachiasmatic nucleus, as well as in retinal amacrine and bipolar cells. The conclusion that Shepherd draws is that significant computing takes place in the dendrites, including in microcircuits in individual dendritic spines. If, as Shepherd suggests, we take the synapse as the basic computational unit of the brain, then we find that the computational resources of the brain are much greater than we would suppose from counting neurons. Since there are about 300 million (300×10^6) synapses per cubic mm of cortex (but only 50 000 neurons), the human brain comprises some 60 trillion (60×10^{12})

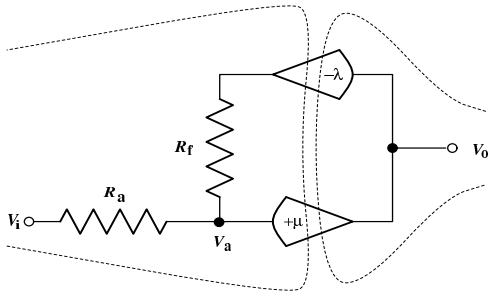

Figure 1: A reciprocal synapse between two dendrites provides negative feedback, in effect implementing an operational amplifier, one of the principal components of an analog computer.

synapses, but only 10 billion neurons (Beaulieu & Colonnier 1983; Shepherd 1990, p. 7).[4]

3.2 Pribram's Neural Wave Equation

Pribram's *holonomic brain theory* (1991) is a development of his earlier holographic hypothesis (1971). Its main points can be summarized as follows: (1) The neural (axonal) net is regular and sparsely connected. That is, in contrast to the full or random interconnection of most PDP models, axon bundles make topology-preserving projections from one area to another. (2) On the other hand, the dendritic net (neuropil) is randomly and densely connected. In fact, if one looks at the neuropil, it appears to be a dense feltwork of dendrites, with cell bodies and axons embedded in this matrix (Fig. 2). Thus PDP models are more similar to the dendritic net than to the axonal net. (3) The foregoing implies that the function of the axons is *communication*, since they make regular projections and impulses are ideally suited to long-distance transmission. (4) On the other hand, the function of the dendrites is *computation*, since they are ideally suited (as we will show below) to subtle, spatiotemporal analog interactions. (5) Indeed, the dendritic net may be viewed as a medium for linear wave interactions. (6) Nonlinearity enters only at axon hillock.

[4]These microcircuits are also sensitive to the geometry of the dendritic tree, since the propagation of the electrical wave is influenced by membrane capacitance, ion conductances, dendritic diameter and many other factors.

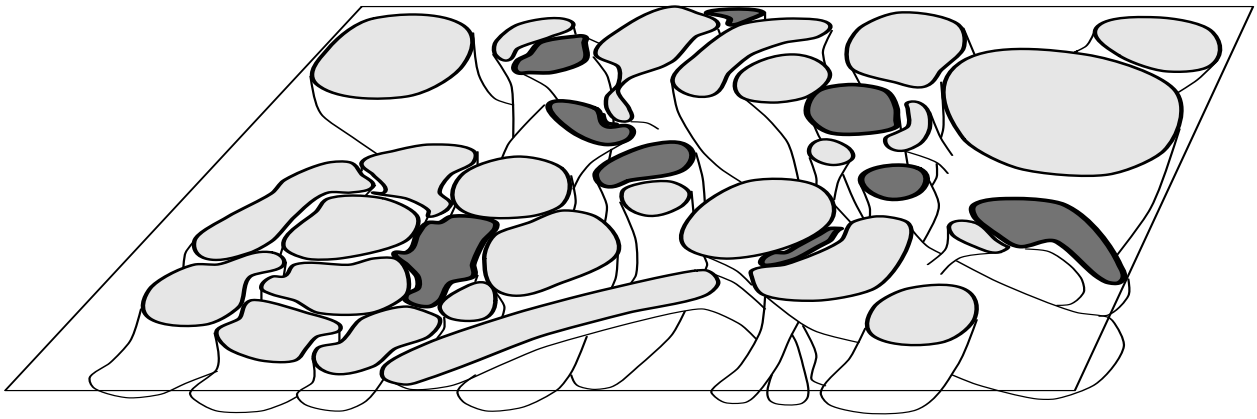

Figure 2: Visualization of neuropil. The figure depicts a cross-section through the neuropil (rat cerebral cortex), magnified approximately 3500 times. Notice the dense packing, allowing for both chemical and electrical interactions. Darker areas represent axons, lighter represent dendrites. Black bands represent synapses (both symmetric and asymmetric). Several dendritic spines are also visible.

3.3 Motivation for Linear Systems Approach

In the Appendices of his 1991 book Pribram, together with Kunio Yasue and Mari Jibu, has presented a neural wave equation for the dynamics of wave interactions in dendritic networks. Several plausible assumptions lead to a wave equation formally identical to that used in quantum mechanics, but this is a two-edged sword: it is very suggestive, and allows the application of many of the mathematical tools developed for quantum theory, but it inherits all the mathematical intractability of quantum mathematics. Analytical results require gross simplifying assumptions, such as a toroidal topology for the dendritic membrane. Many of the results are hard-won, and come as much from the structure of the underlying Hilbert space as from the neural wave equation.

Consideration of these difficulties has led me to try a simplification. By replacing the linear partial differential equations of the neural wave equations by linear ordinary differential equations relating membrane properties at discrete sites, the dendritic net can be viewed as a *dynamic lumped parameter system*. This places it squarely within linear systems theory, an extremely well-understood mathematical framework. Although this approx-

imation misses some of the subtlety of the dendritic interactions, I believe that it will be a more powerful tool for probing the *function* — in contrast to the merely the *dynamics* — of dendritic interactions. Section 6 gives some preliminary results of this kind.

It may be objected that in trying to simplify the linear neural wave equation, I am going in the wrong direction, for there is a widespread opinion, especially in the neural network and neural modeling communities, that linear systems are computationally impotent. For example, Poggio and Torre (1981) say, "information processing by synapses and dendrites must rely on essentially nonlinear interactions between electrotonic signals if it is to be nontrivial," and, "nonlinearities are critical in making any signal transduction interesting from the point of view of information processing." Also Reichardt & Poggio (1981, p. 187) say "every nontrivial computation has to be essentially nonlinear, that is, not representable (even approximately) by linear operators." While I am very sympathetic to the importance of nonlinear systems, I think we may have sold short the information processing capabilities of linear systems. For example, in a forthcoming paper David Wolpert and I show that a linear field computer can simulate a Turing machine (Wolpert & MacLennan submitted), which is hardly trivial and uninteresting information processing.[5]

4 Are Neurons Linear?

4.1 Dendritic Nets

Before we can apply linear systems theory to the dendritic net we must address a factual question: Are synapses linear? Or, more in a accord with the Robustness Principle: Are synapses approximately linear? This question is surprisingly hard to answer from the literature, although Koch & Poggio (1987) is valuable. It is obvious that the synapse cannot behave linearly at high neurotransmitter concentrations since either the receptors become

[5]So that there is no mystery how this is possible, let me simply point out that if the states of a machine are represented by orthogonal vectors, then any deterministic state transition is a simple linear operator (a sum of outer products). The nonlinearities occur in the input coding, not in the computation or input. (This is not, however, the construction used in Wolpert & MacLennan submitted.)

saturated or the neurotransmitter becomes depleted (Koch & Poggio 1983, p. 462). Thus the dependence must become sublinear at the upper end, probably approximately logarithmic.

At low conductances, the postsynaptic potential depends linearly on the synaptic conductance induced by the presynaptic potential (Koch & Poggio 1983, p. 462; Shepherd 1990, p. 427).[6] Although Shepherd claims that we do not know whether synapses operate in their linear range, in simulated dendrites the linear region seems to extend to 10 000 pS (Shepherd 1990, p. 427). Since the quantum of conductance change is on the order of 20–40 pS, and an AP releases less than 200 quanta (8000 pS), it appears likely that synapses will usually be linear. Koch & Poggio (1983, p. 471) suggest that a dendritic spine is in its linear range when the conductance change is less than 10% of the input admittance of the spine head (i.e., a conductance change of less than about 10 000 pS). Furthermore, Koch & Poggio (1985, p. 653) suggest that linear operation can be enhanced by having the input distributed to several sites that are somewhat decoupled in their electrical behavior. Indeed we find that a cortical cell will often make 5–10 synapses with each neuron to which it projects (Shepherd 1990, p. 399), thus improving linearity. Overall we can tentatively conclude that a synapse often behaves linearly, but saturates at high levels. By the Robustness Principle, we can treat the synapse as approximately linear.[7]

[6]Synaptic communication can be decomposed into three components: (1) the relation between presynaptic depolarization and the amount of neurotransmitter released, (2) the relation between quantity of neurotransmitter released and the number of postsynaptic ion channels opened, and (3) the relation between the number of ion channels opened and postsynaptic voltage and current. All three seem to be approximately linear except for saturation. However, (1) can be nonlinear due to facilitation and fatigue (Berne & Levy 1983, p. 60). Nonlinearities enter in (3) because the increased conductance partly depolarizes the membrane, which results in decreased driving potential (Koch & Poggio 1983; Rall & Segev 1987, p. 611; Shepherd 1990, p. 427); for modest depolarizations it is approximately linear (see above).

[7]The chemical synapse operates as voltage-controlled voltage-source or as a voltage-controlled current-source depending on the ratio of the stimulus-induced synaptic conductance to the conductance of the spine neck (Koch & al. 1992). The chemical synapse can also be thought of as a voltage amplifier; a typical gain is 3 (in the linear range), but gains as low as 0.3 and as high as 50 have been observed (Koch & Poggio 1987, pp. 648–649). Chemical synapses can also function as both positive and negative resistances, since presynaptic depolarization can cause postsynaptic hyperpolarization (Poggio & Koch 1985).

4.2 The Axon Hillock

Although my primary concern is with dendritic computation, it will be worthwhile to consider the linearity of the axon hillock. Conventional neural network models are based on the assumption that axonal impulse rate is related to soma depolarization by a sigmoidal or "squashing" function, that is, by a saturating linear function such as $y = \tanh(x)$. However, this model diverges from the biology in several respects (Shepherd 1990, pp. 411–413). First, due to neural adaptation the pulse rate typically decreases during the spike train. For example, if we consider the relation of interspike interval to depolarizing current in certain pyramidal cells, the same depolarizing current will produce a 280/sec. rate for the first interspike interval, but 120/sec. for the second, and only 40/sec. for the fifth. Second, the relation is flatter for later intervals than for earlier. Thus 0.3 nA current leads to a 60/sec. change in rate for the second intervals, but only 40/sec. for the fifth. Finally, the rate/current relation is more linear for later intervals, with the first interval displaying a somewhat irregular (not even monotonic) sigmoidal shape. Although in some cases the axon hillock may function as an approximately linear amplitude to pulse rate converter, for others the relation will be nonlinear. Whether or not a particular neuron's hillock can be treated as linear will therefore depend on its usual operating range.

Shunting inhibition may be a mechanism for nonlinear synaptic interactions, and a basis for multilinear computation, as described in MacLennan (1987, 1990, this volume); for example, Shepherd (1990, p. 417) claims that in effect these inhibitory synapses divide (multiply by a fraction). However, the simulation studies of Koch & al. (1983) suggest that shunting inhibition is more likely to function as a veto of distal excitatory connections; elsewhere they discuss possible multiplicative effects of excitatory and inhibitory synapses on a single spine (Torre & Poggio 1978, 1981), but these effects depend on rather precise timing of signals and placement of synapses.

Nonlinear effects may also enter into the "addition" of inputs from separate spines. However, if the spine neck resistance is high, then the currents are small and combine linearly (Jack & al. 1975, pp. 192–213; Koch & Poggio 1983, pp. 468–469, 1985, pp. 651–653; Shepherd 1990, pp. 465–473). A final source of nonlinearity is the generation of action potentials in the dendrites (Eccles 1957, p. 270; Jack & al. 1975, pp. 213–218; Shepherd 1990, pp. 428–429, 465); although I leave these out of this model, they might be accommodated in the same way as approximately linear axon hillocks (see below).

5 The Linear Systems Model

5.1 Spatiotemporal Fields

Out of the many properties of nervous tissue, we can expect some to be relevant to information processing and others not. Significant quantities likely include membrane depolarizations, ionic currents, transmembrane charges, etc., but may also include others, such as flow rates in microtubules. In contrast with other mathematical models of the detailed dynamics of neural tissue (e.g., Deutsch & Micheli-Tzanakou 1987; Jack & al. 1975; Koch & Poggio 1985), I will not be particularly concerned with the physical quantities that the variables represent. For my purposes, it is not so important to have a final list of the relevant variables; it is important only that they vary approximately continuously (cf. Complementarity) and have approximately linear dynamics.[8]

For mathematical convenience, I assume a discrete set of variables, $\psi(k,t) = \psi_k(t)$, for k in $1,\ldots,N$. These could be, for example, depolarizations at synapses. However, the number of synapses is sufficiently large (say, 5000 to 200 000) that I will often view the variables as $\psi(x,t)$ with x varying over a continuum D, that is, as spatially-continuous time-varying fields, $\psi(t)$.[9] The Complementarity Principle permits this. Further, since we usually take time also to vary continuously, we are really dealing with spatiotemporal fields, $\psi(x,t)$.

Although it is useful to think of the variables as spatiotemporal fields, it is critical that we treat them mathematically as discrete variables, since

[8]The basic theory of information processing presented here would hold for any continuous quantities with approximately linear interactions. For example, this could include certain aspects of the dynamics of the cytoskeleton, as described by Dayhoff and Hameroff (see also Hameroff 1987). In addition to ionic currents in the microtubules, it could include MAP-induced field effects in the average position of vibrating tubulin sidearms. This follows also from the general correspondence between electrical and chemical networks (Oster & al. 1971; Poggio & Koch 1985).

[9]A basic tenet of field computation, as I have defined it (1987, 1990), it that the number of computational units be sufficiently large that it may be treated as a continuous quantity. Thus the Complementarity Principle may be applied. It should be noted that my variety of field computation deals predominantly with nonholonomically constrained fields, as opposed to the holonomically constrained fields familiar from physics (cf. Kugler, these proceedings). However they have in common that they deal with physical quantities varying continuously over a region of space.

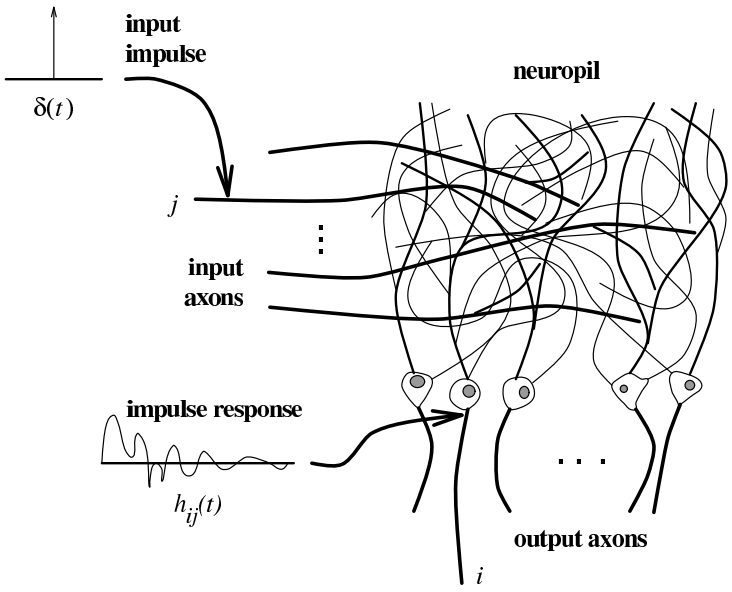

Figure 3: Resonance model of dendritic computation. A dendritic net can be thought of as a large system of coupled resonators with many resonance modes. Input signals arriving on axons activate some of these modes. The resulting signals are integrated in neuron somas and translated into spike trains for transmission to other dendritic nets.

that is the basic simplification over the neural wave equation. Technically, it means that we are viewing the dendritic net as a *lumped parameter system.* In concrete terms, we are viewing the neuropil as a network of discrete summers, multipliers, integrators and differentiators operating on variables of several kinds (e.g., concentrations of different ions). However we can use as fine a division as we choose. Such systems are especially well understood, because they have rational transfer functions.

5.2 Resonance Model

We can think of the overall dynamics of the dendritic processes in terms of a *resonance model* (Fig. 3). The idea is that the dendritic net can be thought of as a large system of coupled resonators with many resonance modes — on the order of the number of synapses (5000 to 200 000), as will be shown below. The input signal drives the net and activates some of these resonances.

The soma, acting analogously to an antenna, integrates the internal signals over space and time. Thus integrated, the activated resonances, if they are sufficiently strong, may trigger action potentials at the axon hillock.

5.3 Dynamic Equation

The *coupling coefficients* of the differential equations are assumed constant. These coefficients include the synaptic efficacies, membrane capacitances, spine-neck diameters, etc. Thus I assume no learning or development during dendritic processing. In effect I analyze the dendritic net at two time scales: the fast scale deals with dendritic resonance, the slow scale with learning. (I will show later an example of how learning is handled.) Next is a mathematical description of the dynamics of dendritic resonance. Readers familiar with linear systems theory will probably want to skip to Section 6 (Examples of Dendritic Information Processing).

A system of higher-order differential equations can be reduced to first order by introducing additional variables ψ_k. Thus the n-th order differential equation

$$\psi^{(n)} = G[t, \psi, \psi^{(1)}, \psi^{(2)}, \ldots, \psi^{(n-1)}]$$

becomes a system of first-order equations:

$$\begin{aligned}
\dot{\psi}_{n-1} &= G(t, \psi, \psi_1, \psi_2, \ldots, \psi_{n-1}) \\
\dot{\psi}_{n-2} &= \psi_{n-1} \\
&\vdots \\
\dot{\psi}_1 &= \psi_2 \\
\dot{\psi} &= \psi_1.
\end{aligned}$$

Henceforth, without loss of generality, I restrict my attention to first-order equations.

The behavior of a dendritic net is determined by three spatiotemporal fields (or sets of variables): (1) the inputs φ_j, normally thought of as depolarizations of the boutons of incoming axons, (2) the outputs ω_i, normally thought of as depolarizations at the axon hillock of outgoing axons, and (3) the state variables ψ_k, which include all properties of the dendritic net relevant to information processing (including those introduced in the reduction to first-order equations). (See Fig. 4.) We will sometimes write the equations

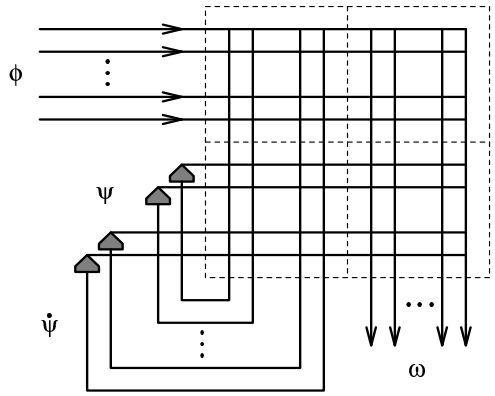

Figure 4: Schematic diagram of dendritic net. Input comes from the left, output leaves at the bottom. Shaded triangular figures are integrators. Intersecting lines within dotted boxes represent connections from horizontal lines to vertical lines with constant efficacies (weights).

in terms of time-varying discrete variables, e.g., $\psi_k(t)$ or $\psi(k,t)$, and other times in terms of time-varying fields, e.g., $\psi(t)$.

The state field $\psi(t)$ is driven linearly by the input field $\varphi(t)$:

$$\dot{\psi}(k,t) = D_{k1}\varphi(1,t) + D_{k2}\varphi(2,t) + \cdots + D_{kn}\varphi(n,t) + \text{other terms},$$

or, more compactly,

$$\dot{\psi}(t) = D\varphi(t) + \text{other terms}.$$

The output field is a sum of the input and state fields:

$$\begin{aligned} \omega(k,t) \;=\; & E_{k1}\varphi(1,t) + E_{k2}\varphi(2,t) + \cdots + E_{kn}\varphi(n,t) + \\ & G_{k1}\psi(1,t) + G_{k2}\psi(2,t) + \cdots + G_{kp}\psi(p,t). \end{aligned}$$

That is:

$$\omega(t) = E\varphi(t) + G\psi(t).$$

The state field is also driven by feedback from itself:

$$\dot{\psi}(k,t) = \text{input drive} + F_{k1}\psi(1,t) + F_{k2}\psi(2,t) + \cdots + F_{kp}\psi(p,t).$$

That is:

$$\dot{\psi}(t) = D\varphi(t) + F\psi(t).$$

The feedback matrix F determines the dynamics of system and hence its resonances.

In summary, the field evolution equations for a dendritic net are:

$$\begin{aligned} \dot{\psi}(t) &= D\varphi(t) + F\psi(t), \\ \omega(t) &= E\varphi(t) + G\psi(t). \end{aligned}$$

By Complementarity, these equations may be interpreted either as system of linear first order ordinary differential equations or as a linear integro-differential equation.

5.4 Laplace Transform Analysis

The system of differential equations can be solved in the usual way by taking the Laplace transform of the evolution equation:

$$\begin{aligned} s\Psi(s) - \psi(0) &= D\Phi(s) + F\Psi(s), \\ &\quad \text{where } \Psi(s) = \mathcal{L}\{\psi(t)\} \\ &\quad \text{and } \Phi(s) = \mathcal{L}\{\varphi(t)\}. \end{aligned}$$

($\mathcal{L}\{\}$ is the Laplace transform.) Define the transformed transition matrix:

$$R(s) = (s\mathrm{I} - F)^{-1},$$

and then the transformed solution is given by:

$$\Psi(s) = R(s)[D\Phi(s) + \psi(0)].$$

For understanding the signal processing properties of the net it is useful to consider its *forced response*, which is its response with no stored energy. Thus we set $\psi(0) = 0$, and the transformed output is:

$$\Omega(s) = E\Phi(s) + GR(s)D\Phi(s).$$

Define the *transfer function matrix*:

$$H(s) = E + GR(s)D,$$

and then the system function in the transformed domain is:

$$\Omega(s) = H(s)\Phi(s).$$

Suppose all the inputs are clamped at zero except the jth input, $\varphi(j,t)$, and consider output i, $\omega(i,t)$. The transforms of this input/output pair determine $H_{ij}(s)$:

$$H_{ij}(s) = \Omega(i,s)/\Phi(j,s).$$

This is the transform of the *impulse response*, $h_{ij}(t)$, which can be convolved with the input to yield the output:

$$\omega(i,t) = h_{ij}(t) * \varphi(j,t).$$

(Throughout this paper "$*$" is the convolution operation.) In particular, if φ_j is a Dirac delta (impulse) function, then:

$$\omega(i,t) = h_{ij}(t).$$

That is, injecting an impulse into input j causes the signal $h_{ij}(t)$ at output i. This is only approximate, since a nerve impulse (an action potential) is a poor approximation to a delta function; nevertheless it suggests the general form of the output (h_{ij} blurred by convolution with the action potential).

To determine the potential number of resonances of a dendritic net, recall that the inverse of a matrix is its adjoint divided by its determinant, and write the transformed transition matrix:

$$R(s) = \frac{\mathrm{adj}(s\mathrm{I} - F)}{|s\mathrm{I} - F|}.$$

There will be poles (infinities) wherever $s\mathrm{I} = F$, that is, at any s that is an eigenvalue of F. Thus the eigenvalues of F are its resonances, that is, the complex frequencies at which the resonances occur. Since F is $p \times p$ matrix, where p is the number of state variables, F has up to p eigenvalues, and so the dendritic net has up to p resonances. But the number of state variables is at least the number of synapses in the dendritic net, which is on the order of 5 000 to 200 000 for each neuron. Therefore I conclude that a typical dendritic net could have 10^4 to 10^6 resonances, which suggests a vast pattern recognition capacity for a dendritic net.

6 Examples of Possible Dendritic Information Processing

6.1 Gabor Coefficients and Dendritic Iteration

The Gabor elementary functions of time (for a given Δt) are defined (Gabor 1946):

$$\gamma_{mn}(t) = \exp\left[-\pi\frac{(t-n\Delta t)^2}{2(\Delta t)^2}\right]\exp\left(\frac{2\pi i m\, t}{\Delta t}\right).$$

There is considerable evidence (summarized in MacLennan 1991) that Gabor functions of space and (perhaps) time are representational primitives in visual cortex. If this is so, then the visual cortex has to find the Gabor coefficients c_{mn} such that:

$$\varphi(t) = \sum_{mn} c_{mn}\gamma_{mn}(t).$$

Unfortunately, the Gabor functions are not orthogonal, and so the coefficients c_{mn} cannot be computed by a simple inner product, $\langle\gamma_{mn},\varphi\rangle$. One way to determine the Gabor coefficients is by performing gradient descent in error of representation (Daugman 1988):

$$E(c) = \left\|\varphi - \sum_{mn} c_{mn}\gamma_{mn}\right\|^2.$$

This can be accomplished by the linear system:

$$\begin{aligned}\dot{c} &= Fc + D\varphi\\ &\quad\text{where } D = \eta G^{\mathrm{T}}\\ &\quad\text{and } F = -\eta G^{\mathrm{T}}G\\ &\qquad\text{where } G = (\ldots,\gamma_{mn},\ldots)\\ &\qquad\text{and } \eta \text{ is the adaptation rate.}\end{aligned}$$

Thus $F_{jk} \propto \langle\gamma_j,\gamma_k\rangle$ for some enumeration γ_j of the Gabor functions γ_{mn}. This system is an iterative correction ($-\eta G^{\mathrm{T}}G$) of the inner product (ηG^{T}). We have seen that such a linear system is a simple computation for a dendritic net. The same approach works for many other nonorthogonal representations.[10]

[10]It should be noted that Gabor (1946) showed that the optimal real, as opposed to complex, basis functions are derivatives of Gaussians, which is the same as saying Gaus-

Iteration is generally considered too slow for biological neural networks, since neuronal signalling limits its speed to 10 to 100 msec./iteration. However, we can see that iteration in dendritic nets is much faster, since the delay of a chemical synapse is only about 0.5 msec. Thus dendritic iteration may be 20 to 200 times as fast as neuronal iteration. Since electrical synapses are essentially delayless, they may mediate even faster computational processes. Thus dendritic computation opens up an entire new range of possible neural information processing algorithms.

6.2 Orthogonality in High-dimensional Spaces

It will become apparent that I will be making considerable use of matched filters, which compute the inner products of patterns, and I will also be considering Hebb-like (i.e. correlational) learning rules. It is widely believed in the neural net community that inner-product pattern matching and Hebbian learning are weak, since nonorthogonal patterns result in excessive crosstalk. Therefore I digress to consider the issue of orthogonality in high-dimensional spaces; recall that the input dimension of a dendritic net is on the order of hundreds of thousands.

I begin with an argument sketched in Hamming (1986), which demonstrates that if we pick randomly any two of the 2^n bipolar vectors in n-dimensional space $\{-1,1\}^n$, then they are almost surely nearly orthogonal. Specifically, with increasing n the cosine of the angle between them approaches 0 almost certainly (by the weak law of large numbers). Alternatively, if we normalize these vectors, then their inner product approaches zero almost certainly:

$$\mathbf{X}^{\mathrm{T}}\mathbf{Y} = \sum_k \pm 1/n,$$

where $\mathbf{X}$ and $\mathbf{Y}$ are random bipolar vectors of the form

$$(\pm 1, \pm 1, \ldots, \pm 1)^{\mathrm{T}}/\sqrt{n}.$$

Since dendritic computation is predominantly analog, we need to consider extensions of Hamming's result to continuous-valued vectors. Let $\mathbf{X}$ and $\mathbf{Y}$

sians times Hermite polynomials or Gaussians times Hermite functions (Stork & Wilson 1990). This set of functions has the added advantage that those defined at a fixed (spatial or temporal) location are orthogonal.

be random vectors in $\mathbf{R}^n$ with zero mean and standard deviation proportional to $1/n$ (thus maintaining normalization on the average). Of course the expectation value of $\mathbf{X}^{\mathrm{T}}\mathbf{Y}$ is 0, but it is also easy to see that its variance is $1/n^3$:

$$\mathrm{Var}\{\mathbf{X}^{\mathrm{T}}\mathbf{Y}\} = \sum \mathcal{E}\{X_k^2\}\mathcal{E}\{Y_k^2\} = \sum \frac{1}{n^2}\,\frac{1}{n^2} = \frac{1}{n^3}.$$

Once again the weak law of large numbers tells us that with increasing n the inner product will almost surely be 0.

We have seen that neural computation and communication is limited to about two digits of precision — most likely even less. Therefore a matched filter in the dendritic net will be unable to distinguish orthogonal vectors from vectors whose inner products are less then ϵ, where ϵ is at least 0.01. Therefore I will consider sets of vectors that are *ϵ-orthogonal*, that is, vectors whose inner product is less than ϵ.[11] For simplicity I assume the vectors are normalized and that $\epsilon \in (0,1)$. Kainen (1992) has shown that for fixed ϵ the number of ϵ-orthogonal n-dimensional vectors increases exponentially in n. In particular, if $N(n)$ is the number of n-dimensional ϵ-orthogonal vectors, then

$$(1-\epsilon^2)^{-1/2} \leq \lim_{n\to\infty} N(n)^{1/n} \leq (1-\epsilon)^{-1/2}.$$

For example, if $\epsilon = 0.1$, then

$$1.00504 \leq \lim_{n\to\infty} N(n)^{1/n} \leq 1.0541.$$

We can see that asymptotically

$$1.0541^n \geq N(n) \geq 1.00504^n,$$

approximately. For example, for $n = 5000$ we have the lower bound $N(5000) \geq 8.2 \times 10^{10}$; for $n = 10\,000$ we have $N(10\,000) \geq 6.7 \times 10^{21}$. So, linear algebra tells us that a 5000-dimensional space has 5000 orthogonal vectors, but we see that it has 82 billion 0.1-orthogonal vectors. By not requiring exact orthogonality we greatly increase our representational resources.[12]

[11] Over the last two decades several investigators have independently invented ϵ-orthogonality; citations can be found in Kainen (1992).

[12] If we pick $\epsilon = 0.01$, the bounds are much less impressive, since we get $8.2 \times 10^{10} \geq N(5000) \geq 1.28$ and $6.7 \times 10^{21} \geq N(10000) \geq 1.65$. It may be that this ϵ is too close to perfect orthogonality. More likely, the problem may simply be that the formula gives very loose bounds. If so, the results for $\epsilon = 0.1$ are even more impressive.

Suppose that inner products less than ϵ are indistinguishable from inner products equal to zero. Then an outer-product associative memory could hold $N(n)$ associated pairs, which we have seen can be much greater than the expected capacity, n pairs. However, this requires the stimulus patterns to be drawn from an ϵ-orthogonal set. This could be accomplished either by a self-organizing process,[13] or by random selection of vectors perhaps followed by a Gram-Schmidt process.

In conclusion, it is well-known that many simple, efficient neural information processes work best on nearly orthogonal representations, since orthogonality decreases crosstalk. But in a space of high dimension, randomly chosen representations will almost always be nearly orthogonal. Thus, in the context of the brain, which is insensitive to slight deviations from orthogonality, nonorthogonality will take care of itself.

6.3 Matched Filters

The inner product $\langle \zeta, \varphi \rangle$ is maximized for identical normalized signals ζ and φ, which makes it a popular way means of pattern matching in neural networks.[14] The inner product can also be expressed as the final value of a reverse convolution (essentially a correlation):

$$\langle \zeta, \varphi \rangle = \zeta(T - t) * \varphi(t)\big|_{t=T}. \tag{1}$$

This is called a *matched filter* for the pattern ζ. Under reasonable assumptions, matched filters are optimal in maximizing signal-to-noise ratio at time T (e.g., see Cooper & McGillem 1986, Sec. 9.5). A matched filter has an impulse response that is the time-reverse of the pattern for which it is tuned:

$$h(t) = \zeta(T - t).$$

From the Eq. 1 for the inner product, we can see that if a dendritic net implements a matched filter, then an action potential will be generated at

[13]The problem of finding an ϵ-orthogonal set is equivalent to the problem of packing s-spheres into a related space (the n-dimensional elliptic space), where $s = (1/2)\arccos\epsilon$ (Kainen 1992). This should be easy so long as we don't try to pack the spheres too tightly.

[14]For notational clarity, these derivations apply to scalar in / scalar out nets, but can be easily extended to fields.

time T if the match exceeds the threshold at the axon hillock. If this threshold is a little less than the norm of ζ (and input φ is approximately normalized), then the matched filter acts as a pattern recognizer for ζ.

There are several ways that we may interpret the output of a matched filter. First, it may be interpreted as the inner product of *spatiotemporal* fields ζ, φ:

$$\omega(T) = \langle \zeta, \varphi \rangle.$$

Second, it may be interpreted as the sum of the pairwise *temporal* correlations of input signals φ_j and stored patterns ζ_j:

$$\omega(T) = \sum_j \langle \zeta_j, \varphi_j \rangle.$$

Finally it may be interpreted as the *temporal* integration of the running *spatial* inner products:

$$\omega(T) = \int_0^T \zeta(t) \cdot \varphi(t) \mathrm{d}t.$$

Each of these interpretations may be informative in different circumstances.

One of the simplest ways of implementing a matched filter is a *tapped-delay line* or *transversal filter*:

$$\omega(t) = \sum_{jk} W_{jk} \zeta_j(t - \tau_k),$$

where the τ_k are delays. This formula is essentially a (discrete) convolution. The self-organization of a matched transversal filter can result from a simple Hebb-like rule. Suppose the firing of ω causes W_{jk} to move toward the last value it transmitted (a mechanism for this is described in Section 6.4):

$$\Delta W_{jk} = c\zeta_j(t - \tau_k).$$

Then each weight approaches the average value that it transmits when ω fires:

$$W_{jk} \longrightarrow \mathcal{E}\{\zeta_j(t - \tau_k) \mid \omega \text{ fires}\}.$$

Thus the dendritic net becomes a filter matched for $\mathcal{E}\{\zeta \mid \omega \text{ fires}\}$, the expected input that causes ω to fire. If there is lateral inhibition among a

number of dendritic nets of this form, then they will tend to become matched filters for distinct spatiotemporal patterns.

Although the transversal filter is simple, it is more efficient (fewer synapses) to make use of the dynamic properties of dendritic feedback. This can be accomplished by an autoregressive implementation of a matched filter:

$$\psi(t) = \sum_i F_i \psi(t - \tau_i) + \sum_j D_j \varphi(t - \tau_i).$$

We can adapt the weights by using the Feintuch (1976) approximate gradient descent algorithm:

$$\begin{aligned} \dot{F}_i &= k[\zeta(T - t) - \psi(t)]\psi(t - \tau_i), \\ \dot{D}_j &= k[\zeta(T - t) - \psi(t)]\varphi(t - \tau_i). \end{aligned}$$

This is a simple "delta rule." An even better approach is to implement the self-organization of the poles and residues of a recursive filter (described in Section 6.6).

Successive layers of self-organizing matched filters will adjust to recognize spatiotemporal patterns at higher levels of abstraction. The higher levels are characterized by longer stored pattern lengths (i.e. larger T in the impulse response $\zeta(T - t)$), which is simply the temporal analog of the successively wider receptive fields found in higher vision areas. Conversely, as shown in the next section, triggering of higher-order patterns will flow backward to generate expectations for lower-order patterns, in effect priming the matched filters.

6.4 Action Potential-Triggered Antidromic Dynamics

An AP (action potential) causes an electric wave to spread electrotonically from the axon hillock back into the dendritic net (Shepherd 1988, p. 137). This *antidromic* electrical signal is transferred efficiently into the dendritic spines (Koch & Poggio 1983, p. 461; Shepherd 1990, p. 464). (See Fig. 5.) To a first approximation this can be treated mathematically as the injection of an impulse (Dirac delta function) at the axon hillock. Since we are putting a signal into the "output" of the dendritic net, we need to know the *antidromic dynamics* of the neuron.

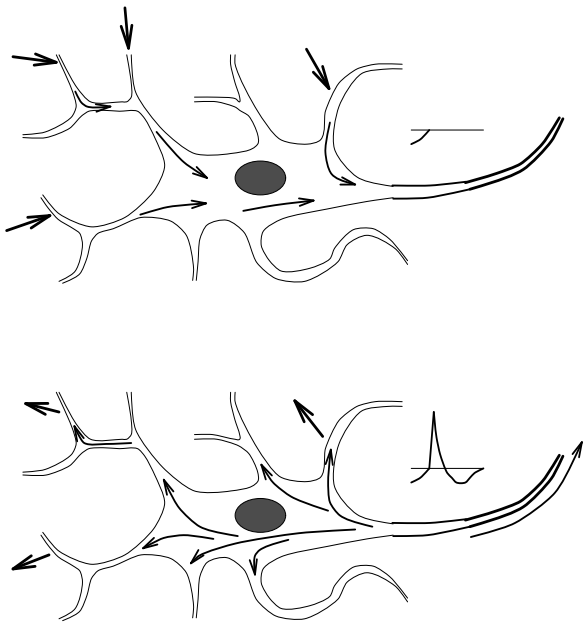

Figure 5: Antidromic electrotonic spread. Converging electric waves from dendritic net are summed in soma. If an action potential is triggered, current is efficiently transferred back into dendritic spines.

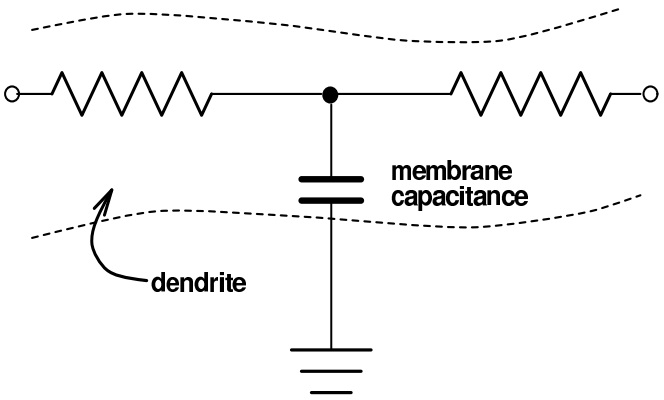

Figure 6: The dendritic membrane capacitance as a reversible (but leaky) integrator. An electrical wave from either direction will charge the capacitance, and so cause the total charge to reflect the (approximate) integral of the signal.

To enable the analysis of the *antidromic dynamics* of the dendritic net, I will assume that the integrators are *reversible* or *reciprocal*. In other words, putting a signal into the output of an integrator causes its integral to appear at the input. This is in fact a reasonable assumption, because integration is usually a result of the membrane capacitance, which may be charged by a electrical wave coming from either direction (Fig. 6). This assumption is also valid for many chemical accumulation processes. For the same reasons we will assume that multipliers (connection strengths) are reciprocal, since they are often a result of electrical resistance, which has the same affect in either direction. Reciprocity (reversibility) does not apply automatically to chemical or electrical synapses, which are usually nonreciprocal or *rectifying* (Koch & Poggio 1987, pp. 649–650), but we ignore this complication for now.[15]

Under the foregoing assumptions it is easy to show that the reversed equations are:

$$\begin{aligned} \varphi &= \psi D + \omega E, \\ \dot{\psi} &= \psi F + \omega G. \end{aligned}$$

[15]It is probably best handled by decomposing F into two matrices, F^{R} and F^{D}, one of which is reversible, the other not. Under some conditions electrical synapses exhibit reciprocity (Koch & Poggio 1987, p. 650).

The transfer function matrix is the transpose of that of the original system:

$$\Phi(s) = H(s)^{\mathrm{T}}\Omega(s).$$

Now we consider the antidromic dynamics in the special case where the dendritic net implements a matched filter (again assuming φ, ω and ζ are scalars). For a matched filter the transfer function is:

$$H(s) = e^{sT}\mathrm{Z}(-s).$$

For the antidromic impulse response, set the output to the transform of a delta function, $\Omega(s) = 1$. Therefore, the antidromic impulse response is:

$$\Phi(s) = e^{sT}\mathrm{Z}(-s).$$

In other words, the signal at the "input" is exactly the pattern to which the matched filter is tuned:

$$\varphi(t) = \zeta(T - t).$$

We have seen that for a matched filter, the pattern for which it is tuned is produced at the dendritic terminals. There are several effects this could have. First, it could trigger or enhance learning (e.g., by a Hebbian rule) by delivering extra current to the activated dendrites. (I showed above how this could function in the self-organization of transversal filters.) Second, the antidromic spread could generate an expectation for the stored pattern by partially depolarizing the dendritic spines, in effect "priming" the dendrites. Such priming could proceed top-down through many levels of successively less abstract matched filters, thus providing a mechanism for top-down expectations. Finally, the antidromic electrotonic flow may implement a kind of pattern completion, since if a partial pattern succeeds in triggering an AP, it will succeed in regenerating the complete pattern at the input terminals. (However, these theoretical predictions have not been tested by simulation using realistic APs.)

In summary, the generation of an AP at the output of a dendritic net will cause production of that output's impulse response at the input terminals of the net. This could be a mechanism for triggering learning, for generating expectations, and for pattern completion. It assumes that (1) integrators and multipliers are reciprocal; (2) the AP can be modeled as a Dirac delta function.

6.5 Spectral Density Analysis

If, as indicated in Berger, Pribram et al. (1990, 1992), the statistical properties of the impulse train are more significant than the actual phase relations, then the spectral density may be more relevant than the Laplace transform as a description of impulse spike trains. Their investigation shows that the results of dendritic computation are encoded in the statistics of spike trains. Specifically, the results determine the barrier height and drift coefficient of a random walk that determines the mean and standard deviation of the interspike intervals. This suggests that spike trains should be treated as random processes, so I turn to the spectral density, which is a powerful tool for such analysis.

The *autocorrelation* of a signal $\varphi(t)$ describes its temporal periodic structure, and is defined:

$$R_\varphi(t) = \varphi(t) * \varphi(-t).$$

The autocorrelation of white noise is a delta function, since it has no periodic structure. The *spectral density* of a signal ω is simply the Fourier transform of its autocorrelation:

$$S_\varphi(s) = \mathcal{F}\{R_\varphi(t)\} = \mathcal{F}\{\varphi(t) * \varphi(-t)\}.$$

Thus the spectral density of white noise is a constant. By the convolution theorem we can express the spectral density as the product of the direct and reversed Fourier transforms:

$$S_\varphi(s) = \Phi(s)\Phi(-s).$$

The spectral density of the output of a linear system is the product of the spectral density of its input and a spectral density transfer function,

$$S_\omega(s) = S_h(s)S_\varphi(s),$$

which is given by $S_h(s) = H(s)H(-s)$, the spectral density of the impulse response of the system. Thus we see that the spectral density transfer function of a matched filter,

$$S_h(s) = \mathcal{F}\{\zeta(-t) * \zeta(t)\} = S_\zeta(s),$$

is simply the spectral density of the pattern to which it is tuned. More accurately, the spectral density is given by $\mathcal{F}\{\zeta(T-t) * \zeta(t-T)\}$, but the phase is ignored by the spectral density:

$$\mathcal{F}\{\zeta(T-t) * \zeta(t-T)\} = e^{sT}\mathrm{Z}(-s)\ e^{-sT}\mathrm{Z}(s) = \mathrm{Z}(-s)\mathrm{Z}(s).$$

As we've seen, the autocorrelation of white noise $W(t)$ is an impulse function, $R_W(t) = c\delta(t)$, and its spectral density is thus a constant, $S_W(s) = c$. Therefore the spectral density of the output of a matched filter that is given white noise is:

$$S_y(s) = S_h(s)S_W(s) = S_\zeta(s)c.$$

Thus the white-noise response is proportional to the spectral density of the pattern to which the matched filter is tuned. In other words, we can discover a matched filter's "expectation" by probing it with white noise.

We've seen that white noise input to a matched filter causes it to reproduce a signal with the same spectral density as the pattern to which it is tuned. I mention briefly the possible significance of this for understanding certain perceptual phenomena. It is well known that experimental subjects in sensory deprivation tanks often experience hallucinations. Further, people often report hearing "radio or TV sounds" in the white noise of forced-air heaters, air conditioners, etc., but cannot identify precisely what they are hearing. In these situations white noise may be showing us the signals to which sensory systems are tuned.

It's noteworthy that many divinatory practices, ancient and modern, rely on: (1) preventing sensory input (covering ears, closing eyes), thus causing the automatic gain control to increase sensitivity; then (2) providing white noise input (rustle of leaves or water, voices of a crowd, reflective or blank surface, etc.); then (3) looking or listening for something relevant. By activating higher-order matched filters with white noise, these practices could expose top-down expectations, concerns, etc. in the mind of the practitioner.

6.6 Self-organization of Recursive Matched Filters

I suggested above that a more efficient use of dendritic processing resources uses recursive (as opposed to transversal) filters, but now we need to consider their self-organization by biologically plausible mechanisms. Suppose

we want to construct a matched filter for ζ of the form:

$$\begin{aligned} \dot{\psi} &= F\psi + D\varphi, \\ \omega &= G\psi. \end{aligned}$$

where G is a row vector and D is a column vector. (For simplicity I will restrict my attention to filters with scalar input and output.) I will assume that there are enough state variables ψ_k to match ζ reasonably well (e.g., at least twice as many as the number of poles of ζ, so that we can have a second order differential equation for each pole).

There are several ways a filter of this kind can self-organize, but the simplest is based on the generalized Fourier series. To have a matched filter for the signal ζ, complete with its phase relations, it's necessary that the impulse response be $h(t) = \zeta(T - t)$, where T is the duration of pattern ζ. Therefore, suppose that ρ_k is an orthonormal system over $(0, T)$ and that $h(t) = \zeta(T - t)$ has the generalized Fourier series:

$$\zeta(T - t) = \sum_{k=0}^{\infty} c_k \rho_k(t),$$

where

$$c_k = \int_0^T \zeta(T - t)\rho_k(t)\mathrm{d}t. \tag{2}$$

The goal is to extract the Fourier coefficients c_k of $\zeta(T - t)$ from a training presentation of ζ to the network. Now observe that by Eq. 2 the coefficients are given by the values at $t = T$ of convolutions:

$$c_k = \zeta(t) * \rho_k(t)|_{t=T}\,.$$

Therefore, if the training signal is presented to filters with impulse responses $h_k(t) = \rho_k(t)$, then their outputs at time T will be the Fourier coefficients. If these coefficients are used as the weights on the filters, then the total impulse response will be

$$h(t) = \sum_{k=0}^{n} c_k h_k(t) = \zeta(T - t),$$

as desired.

This transfer of the Fourier coefficients into the weights could be triggered by an antidromic electrotonic pulse, for example (Fig. 7). More specifically,

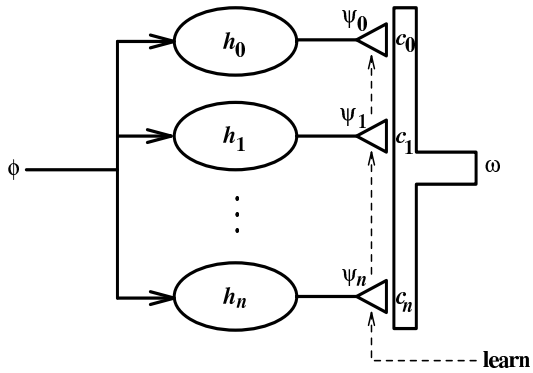

Figure 7: Adaptive recursive filter. In the simplest form the filters h_k have impulse response equal to orthonormal basis functions ρ_k. At time T after a signal ζ is input, the output activities of these filters, ψ_k, are the Fourier coefficients of $\zeta(T-t)$. A learning trigger, for example, an antidromic impulse from the postsynaptic neuron or a learning signal from a third source, can cause these ψ_k to become the synaptic efficacies c_k, thus matching the filter to ζ.

at time T the presynaptic activity, as reflected in depolarization or neurotransmitter release, for example, is proportional to the corresponding Fourier coefficient. All that is necessary to match the network to the signal ζ received up to that point is to cause that activity to determine the synaptic efficacy. Although all the information comes from the presynaptic side, the trigger to cause the efficacy change could be an antidromic electrotonic pulse on the postsynaptic side, or indeed a signal from a third location. Thus the overall scheme is compatible with various kinds of Hebbian learning.

Next consider the dendritic implementation of the filters themselves. The most familiar basis fields ρ_n are the trigonometric functions:

$$\begin{aligned} \rho_0(t) &= 1, \\ \rho_{2k}(t) &= \cos(2\pi kt/T), \\ \rho_{2k-1}(t) &= \sin(2\pi kt/T). \end{aligned}$$

These filters are implemented by simple second order equations:

$$\begin{aligned} \ddot{\rho}_{2k} &= \dot{\varphi} - \nu_{2k}^2 \rho_{2k}, \\ \dot{\rho}_{2k-1} &= \varphi - \nu_{2k-1}^2 \int \rho_{2k-1} \mathrm{d}t, \end{aligned}$$

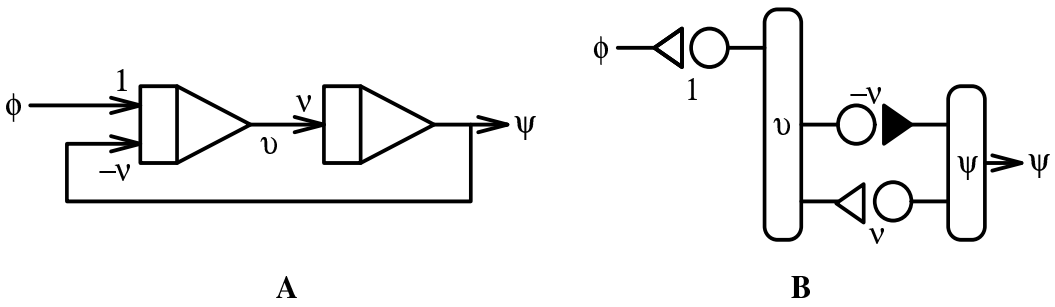

Figure 8: Possible neural implementation of sine filter. **(A)** Typical analog circuit implementation of sine filter; triangular objects are integrators, whose inputs are summed with the indicated weights. **(B)** Possible dendritic implementation of the analog circuit shown in (A). Graded potentials are input on dendritic spines and are summed and temporally integrated in the dendritic trunks or the soma. Integrated sums are transferred into presynaptic dendrites or axons. Open triangles are excitatory synapses; filled triangles are inhibitory.

where $\nu_k = 2\pi k/T$. These filters have a comparatively simple implementation in dendritic circuits; nonlinear saturation effects ensure stability (Figs. 8, 9). The sine filters ρ_{2k-1} can be dispensed with in many cases, since the cosine filters ρ_{2k} are sufficient to recover the spectral density of the signal without its phase, which, as we've seen, may be the usual case.

The regular patterns of weights in the $\boldsymbol{\nu}$ vector may seem biologically implausible. However, if $\boldsymbol{\nu}$ is a uniformly random vector over the same range, the results are not much worse, in spite of there being no guarantee that the filters are orthogonal.

7 Conclusions

In this paper I have argued that (1) mathematical models are at best approximately realized in nervous systems, and so we should take them seriously, but not *too* seriously; (2) linear dynamical systems are useful models of information processing in dendritic nets; (3) iterative processing is feasible in dendritic nets, and an example is the computation of Gabor coefficients; (4) dendritic nets can easily self-organize into matched filters for spatiotemporal fields; (5) white noise may stimulate recall of stored patterns; and (6) impulse-triggered reverse electrotonic spread may facilitate learning and gen-

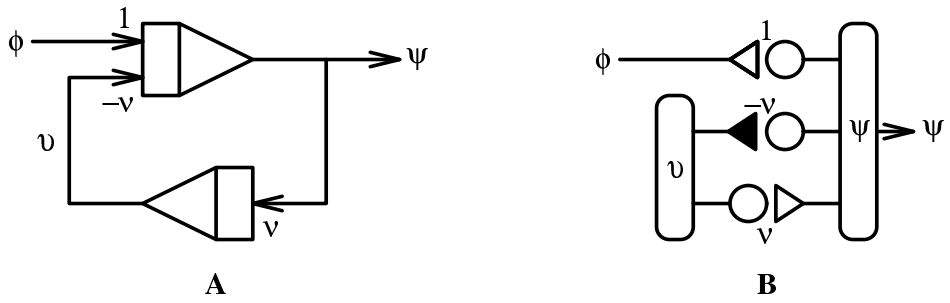

Figure 9: Possible neural implementation of cosine filter. **(A)** Typical analog circuit implementation of cosine filter. **(B)** Possible dendritic implementation of the analog circuit shown in (A).

erate top-down expectations. Overall we have seen that the dendritic net, although approximately linear, may be a powerful processor of information.

8 References

Beaulieu, C., & Colonnier, M. (1983). The number of neurons in the different laminae of the binocular and monocular regions of area 17 in the cat. *J. Comp. Neurol*, **217**: 337–344.

Berger, D., Pribram, K., Wild, H., & Bridges, C. (1990). An analysis of neural spike-train distributions: determinants of the response of visual cortex neurons to changes in orientation and spatial frequency. *Experimental Brain Research*, **80**: 129–134.

Berger, D. H., & Pribram, K. H. (1992). The relationship between the Gabor elementary function and a stochastic model of the inter-spike interval distribution in the responses of visual cortex neurons. *Biological Cybernetics*, **67**: 191–194.

Berne, R. M., & Levy, M. N. (Eds.) (1983). *Physiology.* St. Louis MO: C. V. Mosby.

Cooper, G. R., & McGillem, C. D. (1986). *Probabilistic methods of signal and system analysis*, 2nd ed. New York NY: Holt, Rinehart and Winston.

Crick, F. H. C., & Asanuma, C. (1986). Certain aspects of the anatomy and physiology of the cerebral cortex. In: J. L. McClelland & D. E. Rumelhart (Eds.), *Parallel Distributed Processing: Explorations in the Microstructure of Cognition*, Vol. 2, *Psychological and Biological Models*. Cambridge MA: MIT Press.

Daugman, J. G. (1988). Complete discrete 2-D Gabor transforms by neural networks for image analysis and compression. *IEEE Transactions on Acoustics, Speech and Signal Processing*, **36**: 1169–1179.

Deutsch, S., & Micheli-Tzanakon, E. (1987). *Neuroelectric systems*. New York NY: New York University Press.

Eccles, J. C. (1957). *The physiology of nerve cells*. Baltimore MD: Johns Hopkins Press.

Feintuch, P. L. (1976). An adaptive recursive LMS filter. *Proceedings of the IEEE*, **64**: 1622–1624.

Gabor, D. (1946). Theory of communication. *Journal of the Institution of Electrical Engineers, III*, **93**: 429–457.

Hameroff, S. R. (1987). *Ultimate computing: Biomolecular consciousness and nanotechnology*. Amsterdam: North Holland.

Hamming, R. W. (1986). *Coding and information theory*. Englewood Cliffs NJ: Prentice-Hall.

Jack, J. J. B., Noble, D., & Tsien, R. W. (1975). *Electric current flow in excitable cells*. Oxford UK: Clarendon Press.

Kainen, P. C. (1992). Orthogonal dimension and tolerance. Unpublished report, Washington DC: Industrial Math.

Koch, C., & Poggio, T. (1983). A theoretical analysis of electrical properties of spines. *Proceedings Royal Society London B*, **218**: 455–477.

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

Koch, C., Poggio, T., & Torre, V. (1983). Nonlinear interactions in a dendritic tree: Localization, timing, and role in information processing. *Proceedings National Academy Sciences USA*, **80**: 2799–2802.

Koch, C., Zador, A., & Brown, T. H. (1992). Dendritic spines: Convergence of theory and experiment. *Science*, **256**: 973–974.

MacLennan, B. J. (1987). Technology-independent design of neurocomputers: The universal field computer. In M. Caudill & C. Butler (Eds.), *Proceedings, IEEE First International Conference on Neural Networks* (Vol. 3, pp. 39–49). New York NY: Institute of Electrical and Electronic Engineers.

MacLennan, B. J. (1990). *Field computation: A theoretical framework for massively parallel analog computation; parts I – IV* (report CS-90-100). Knoxville TN: University of Tennessee, Computer Science Department.

MacLennan, B. J. (1991). *Gabor representations of spatiotemporal visual images* (Report No. CS-91-144). Knoxville TN: University of Tennessee, Knoxville, Computer Science Department; submitted for publication.

MacLennan, B. J. (in press-a). Characteristics of connectionist knowledge representation. *Information Sciences*, to appear.

MacLennan, B. J. (in press-b). Continuous symbol systems: The logic of connectionism. In D. S. Levine and M. Aparicio IV (Eds.), *Neural Networks for Knowledge Representation and Inference*. Hillsdale NJ: Lawrence Erlbaum.

MacLennan, B. J. (this volume). Field computation in the brain. This volume.

MacLennan, B. J., & Pribram, K. H. (in preparation). Neural computation without sigmoids.

Neher, E. (1992). Ion channels for communication between and within cells. *Science*, **256**: 498–502.

Oster, G., Perelson, A., & Katchalsky, A. (1971). Network thermodynamics. *Nature*, **234**: 393–399.

Poggio, T., & Koch, C. (1985). Ill-posed problems in early vision: From computational theory to analog networks. *Proceedings Royal Society London B*, **226**: 303–323.

Poggio, T., & Torre, V. (1981). A theory of synaptic interactions. In W. E. Reichardt & T. Poggio (Eds.), *Theoretical approaches in neurobiology* (pp. 28–38). Cambridge MA: MIT Press.

Poggio, T., Torre, V., & Koch, C. (1985). Computational vision and regularization theory. *Nature*, **317**: 314–319.

Pribram, K. H. (1971). *Languages of the brain: Experimental paradoxes and principles in neuropsychology.* Englewood Cliffs NJ: Prentice-Hall.

Pribram, K. H. (1991). *Brain and perception: Holonomy and structure in figural processing.* Hillsdale NJ: Lawrence Erlbaum.

Rall, W., & Segev, I. (1987). Functional possibilities for synapses on dendrites and dendritic spines. In G. M. Edelman, W. E. Gall & W. M. Cowan (Eds.), *Synaptic Function* (pp. 605–636). New York NY: Wiley.

Reichardt, W. E., & Poggio, T., (Eds.). (1981). *Theoretical approaches in neurobiology.* Cambridge MA: MIT Press.

Roberts, A., and Bush, B. M. H. (1981). *Neurones without impulses: their significance for vertebrate and invertebrate nervous systems.* Cambridge UK: Cambridge University Press.

Sakmann, B. (1992). Elementary steps in synaptic transmission revealed by currents through single ion channels. *Science*, **256**: 503–512.

Shepherd, G. M. (1972). The neuron doctrine: A revision of functional concepts. *Yale Journal of Biology and Medicine*, **45**: 584–599.

Shepherd, G. M. (1978). Microcircuits in the nervous system. *Scientific American*, **238**: 92–103.

Shepherd, G. M. (1988). *Neurobiology*, second edition. New York NY: Oxford University Press.

Shepherd, G. M. (Ed.). (1990). *The Synaptic organization of the brain*, third edition. New York NY: Oxford University Press.

Shepherd, G. M. (1990a). The significance of real neuron architectures for neural network simulations. In E. L. Schwartz (Ed.), *Computational Neuroscience* (pp. 82–96). Cambridge MA: MIT Press.

Stork, D. G., & Wilson, H. R. (1990). Do Gabor functions provide appropriate descriptions of visual cortical receptive fields? *Journal of the Optical Society of America A*, **7**, 8 (August 1990): 1362–1373.

Torre, V., & Poggio, T. (1978). A synaptic mechanism possibly underlying directional selectivity to motion. *Proceedings Royal Society London B*, **202**: 409–416.

Torre, V., & Poggio, T. (1981). An application: A synaptic mechanism possibly underlying motion detection. In W. E. Reichardt & T. Poggio (Eds.), *Theoretical approaches in neurobiology* (pp. 39–46). Cambridge MA: MIT Press.

Wolpert, D. H., & MacLennan, B. J. (submitted). A computationally universal field computer which is purely linear.