Field Computation in the Brain*

Bruce MacLennan†

CS-92-174

Abstract

We begin with a brief consideration of the *topology of knowledge*. It has traditionally been assumed that true knowledge must be represented by discrete symbol structures, but recent research in psychology, philosophy and computer science has shown the fundamental importance of *subsymbolic* information processing, in which knowledge is represented in terms of very large numbers — or even continua — of *microfeatures*. We believe that this sets the stage for a fundamentally new theory of knowledge, and we sketch a theory of continuous information representation and processing. Next we consider *field computation*, a kind of continuous information processing that emphasizes spatially continuous *fields* of information. This is a reasonable approximation for macroscopic areas of cortex and provides a convenient mathematical framework for studying information processing at this level. We apply it also to a linear-systems model of dendritic information processing. We consider examples from the visual cortex, including Gabor and wavelet representations, and outline field-based theories of sensorimotor intentions and of model-based deduction.

*Presented at the 1st Appalachian Conference on Behavioral Neurodynamics: Processing in Biological Neural Networks, in conjunction with the Inaugural Ceremonies for the Center for Brain Research and Informational Sciences, Radford University, Radford VA, September 17–20, 1992.

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

1 Topology of Knowledge

1.1 The Assumption that Knowledge is Discrete

For the ancient Greeks the *knowable* and the *sayable* were nearly identical. Socrates is quite explicit: "what we know we must surely be able to tell" (*Laches* 190c), but the idea goes much further back and is nearly inherent in the Greek language: *logos* simultaneously means word, language, thought, reason, explanation, calculation and meaning. Socrates likely came out of the Pythagorean tradition, which reduced the universe to numbers (especially ratios, another meaning of *logos*), and which *calculated* by the mechanical manipulation of formal arrangements of pebbles (*calculi*, in Latin). Through Plato and Aristotle this led to the idea that true knowledge could be reduced to a deductive structure in which inference is represented by the mechanical rearrangement of patterns of discrete symbols. The search for a calculus for knowledge representation and inference was continued by such figures as Lull, Hobbes, Leibniz and Boole.

These efforts reached a kind of culmination in the twentieth century. The development of practical symbolic logics by Peano, Russell, Whitehead and others created the real possibility of putting knowledge in the form of a calculus. Within just a few decades, however, inherent theoretical limitations of calculi were discovered by Gödel, Turing, Löwenheim, Skolem and others. In order to do this they investigated the mechanical manipulation of discrete symbols, and so laid the foundation for the theory of digital computation.

The twentieth century also brought the technological means — the electronic digital computer — for manipulating large discrete-symbol structures at high speed. The theory that knowledge can and must be reducible to calculi then found its home in AI (artificial intelligence), which attempted to apply it to practical problems. The emerging discipline of cognitive science also adopted this view of knowledge in its information processing model of cognition. The assumption that knowledge representation and processing is equivalent to the formal manipulation of discrete symbols was accepted, almost without question, until the mid-1980s, when finally its limitations, both as a technology and a model of cognition, could no longer be ignored.

There is no need to rehearse here the arguments in favor of *connectionist* knowledge representation over the traditional, symbolic approaches. I will observe only that connectionism brings with it a recognition of the role of flexible, context-sensitive information processing as a foundation upon which rest the more symbolic processes. Thus discrete, or approximately discrete, symbol manipulation is viewed as an emergent phenomenon grounded in continuous, or approximately continuous, subsymbolic processes.

We believe that, although there have been many demonstrations of the power of connectionist knowledge representation, its progress is impeded by the lack theoretical construct that captures the essence of connectionist knowledge representation. We have attempted to fill this gap by developing the idea of a *simulacrum*, a model of continuous information representation and processing that fills a role in connectionist epistemology analogous to that filled by the idea of a *calculus* in traditional epistemology (MacLennan in press-a, in press-b, subm.-c). It attempts to answer the question, "What, if not symbols, can be a medium for knowledge representation and processing?" The simulacrum is postulated as the central concept of the theory of continuous computation.

1.2 Theory of Continuous Computation

It must be remarked that a simulacrum, like a calculus, is an idealization of reality. Just as a calculus is taken to be perfectly (i.e. topologically) discrete, so a simulacrum is taken to be perfectly continuous (the mathematical formalization is below). For example, in the conventional theory of (discrete) computation, we assume certain processes are unproblematic, such as the separation of a token from the background or the classification of a token as to its type. Thus, we don't consider the possibility that a Turing machine could misidentify the tape symbol under its read head, although this would be a significant issue for a real (vs. ideal) Turing machine. Similarly, in the theory of continuous computation we assume the continuity of the spaces, maps and processes, even though in practice they might be represented in terms of discrete charge carriers, for example. For both calculi and simulacra the relevant question is whether the real system is sufficiently close to the ideal that the differences may be ignored. In the following we present the simulacrum as an idealized model of continuous computation.

The central idea in the theory of simulacra is the *image*, which is the vehicle of continuous information representation; images correspond to the symbols, formulas and other structures of calculi. The images in a simulacrum belong to one or more *image spaces*, which determine their topology.

Examples of images include the set of all visual images (of bounded area and amplitude) and the set of all auditory images (likewise bounded). On the other hand, a single real number can be considered an image, and an interval of the real line is perhaps the simplest image space. Image spaces satisfy the following postulates.[1]

Postulate 1 *Image spaces are path-connected metric spaces.*

Some implications of this postulate are (1) that images have quantifiable degrees of similarity, (2) that any image is reachable from any other in the space by a continuous process of transformation, and (3) that image spaces have at least the cardinality of the real numbers.

Postulate 2 *Image spaces are separable and complete.*

This postulate is introduced mainly for mathematical convenience; it ensures that images can be approximated by convergent sequences and that the limits of these sequences are in the space. One important consequence of this postulate is that image spaces are topologically equivalent to subsets of Hilbert spaces, which allows us to apply the theory of field computation (see below).[2]

Postulate 3 *Maps between image spaces are continuous.*

One implication of this is that syntactic relations between images are continuous and inherently fuzzy.

Just as in idealized discrete computers the state transitions are taken to occur at discrete time intervals (even though in fact the underlying physical processes proceed continuously), so likewise in idealized continuous computers, states are taken to change continuously (even though some implementations might approximate this with small discrete steps). We define a *formal process* to be one that depends only on the form of the image representing the state of the system; the process is not affected by any meaning that may be associated with the images.

[1] These postulates are tentative; our familiarity with simulacra is too slight to permit a definitive formalization. Justification for the postulates is in MacLennan (in press-a, in press-b).

[2] This follows from a theorem of Urysohn which shows that any metric space with a countable base is homeomorphic to a subset of $\mathcal{L}_2$ (Nemytskii & Stepanov 1989, p. 324).

Postulate 4 *Formal processes in simulacra are continuous functions of time and process state.*

The preceding postulates deal with simulacra as *formal* systems, that is idealized computational processes that depend on the *form* of images but not on any interpretation of them; they are the postulates of *uninterpreted simulacra* or *continuous formal systems*. Now we turn to *interpreted simulacra*, which can be considered *continuous symbol systems*. Thus as the interpretation of a calculus is required to be *systematic*, in particular, to be *compositional*, to respect the constituent structure of the formulas, so also we require systematicity of the interpretations of simulacra, in particular, that interpretations be continuous.[3] Thus:

Postulate 5 *Interpretations of simulacra are continuous.*

There are many open questions in the theory of continuous computation. One immediate question is whether the famous undecidability and uncomputability results of Gödel and Turing apply. Interestingly, some of these questions cannot even be asked in a consistently continuous context, so the problems must be reformulated.[4] Another issue is the existence of *universal machines* for continuous computation. Although it's well-known that under various idealizing assumptions artificial neural networks can simulate Turing machines, and Wolpert & MacLennan (subm.) present a purely-linear continuous-computer with Turing power, we think that there are probably other notions of computational universality that are more appropriate to the theory of continuous computation (MacLennan 1987a).[5]

Another open problem — in continuous computation as well as discrete — is how representations can come to have meanings. Of course, human beings can impose interpretations on otherwise uninterpreted computational systems. But a central philosophical question for computational theories of

[3]Indeed, systematicity in both cases is equivalent to continuity, since respect for constituent structure is just continuity under the appropriate topology for discrete, hierarchically structured formulas (Scott 1970, 1971, 1973; Scott & Strachey 1971).

[4]See MacLennan (in press-a, in press-b) for a discussion and preliminary results. Other formulations of the problems can be found in Blum & al. (1989, 1988), Pour-El & Richards (1979, 1981, 1982) and Stannett (1990).

[5]See also McCulloch & Pitts (1943), Pollack (1987), Hartley & Szu (1987), Franklin & Garzon (1990), and Garzon & Franklin (1989, 1990) for discussion of computational universality in a continuous context.

cognition is whether representations can acquire meanings on their own, so-called *original intentionality*, as opposed to having meanings attributed by a outside observer, *derived intentionality* (e.g., Dennett 1987, 1988). Harnad (1990, in press) has called this the *symbol grounding problem.* Although he thinks it applies only to digital computers and not to analog computers, elsewhere (e.g. MacLennan subm.-c) we argue that grounding is just as much of an issue for continuous (analog) computers as for discrete (digital) computers. As should be clear from the preceding discussion, the constraints of systematic interpretation are no less for continuous computational systems than for discrete, and the solution of the symbol grounding problem does not hinge on the continuous/discrete distinction.[6]

Finally, one of the most important open problems, from the standpoints of both psychology and artificial intelligence, is to understand the emergence of quasidiscrete symbolic processes from the underlying, continuous subsymbolic processes (MacLennan 1992b, in press-a, in press-b). This is the dual problem of that which traditional, symbolic AI was unable to solve: the reduction of continuous information representation and processing (including tacit knowledge, perceptual understanding, sensorimotor skill and associative memory) to discrete symbol manipulation. It resulted in the *cognitive inversion* of the "old AI" (MacLennan 1987b, 1988): it was most successful where humans are least successful (e.g., formal deduction), but least successful where people — and even lower animals — are most successful (e.g., pattern recognition). However, the goal of the "new AI" should not be limited to a connectionist implementation of traditional discrete symbol manipulation (a neural network implementation of LISP); rather we term the symbol processing of the new AI *quasidiscrete* because it has an ineluctable admixture of the continuous, which imparts to it the flexibility characteristic of human symbol use.

[6]We have proposed a different solution to the problem of original intentionality, in which conventional representations acquire meaning through shared relevance to a community that is ultimately grounded in inclusive fitness (MacLennan 1992a; MacLennan & Burghardt subm.). Harnad's notion of grounding would be derivative from this.

2 Field Computation

2.1 Overview of Theory

Evidence is accumulating that Hilbert spaces provide a central theoretical framework in which to construct a theory of neurodynamics and cognition. In addition to the neural wave equation developed by Pribram, Yasue and Jibu (Pribram 1991), we have Urysohn's theorem, cited above, which implies that image spaces are topologically equivalent to subsets of Hilbert spaces. We also expect Hilbert spaces to provide the basis for understanding the emergence of symbolic cognition from subsymbolic processes, since the continuous and the discrete meet mathematically in Hilbert space ($\mathcal{L}_2$ is homeomorphic to ℓ_2); recall that the square wave is an infinite superposition of sinusoids. Finally we have found Hilbert spaces to be the most convenient theoretical framework in which to construct our theory of *field computation*, to which we now turn.

It is necessary to remark on the meaning of *field* as used in the phrase 'field computation' (MacLennan 1987a; MacLennan 1990), which is somewhat broader than current usage in physics, but corresponds to Faraday's. Informally, we define a *field* to be a *spatially extended continuum of quantity.* This seems to be an especially useful basis for understanding the neurodynamics of cognitive processes, as recognized by the Gestalt psychologists (e.g. Köhler 1940, Ch. II). We'll illustrate the idea with several examples before giving a formal definition.

Perhaps the simplest example of a field is the distribution of light intensity over the rods in the retina. Although we know the number of rods is finite, it is sufficiently large (10^8) that they may be treated as a continuum and analyzed through the theory of field computation (MacLennan 1987a). The value of the field φ at retinal position $\mathbf{p}$, $\varphi(\mathbf{p})$, is a scalar representing the light intensity.

A more complex example is provided by the activity of the cones, since there are three kinds of these, and so the activity of the cones is represented by a vector field $\boldsymbol{\varphi}$. The vector $(x, y, z) = \boldsymbol{\varphi}(\mathbf{p})$ represents the activities (x, y and z) of the three kinds of cones at location $\mathbf{p}$.

The preceding examples were two-dimensional fields, that is, quantities defined over a two-dimensional continuum (the retina). One dimensional fields are also common in the nervous system; for example, the instantaneous activities of hair cells in the cochlea define a one-dimensional field $\varphi(f)$, where f is the frequency of the sound. We consider later fields of dimension greater than two.

It's obvious that the fields we've mentioned are *time-varying*, and we have argued elsewhere (MacLennan 1991) for the importance of their temporal structure. Therefore, we often find it convenient to view a time-varying spatially-extended field $\varphi(\mathbf{x})$ as a *spatiotemporally-extended field* $\varphi(\mathbf{x}, t)$. This often yields a considerable theoretical simplification and greater insight into the neural processes, as we'll show later.

Although we've illustrated fields with examples from sensory systems, they are just as prevalent in motor systems and in higher cognitive areas, as will also become apparent later.

Now we define fields more formally. Readers uninterested in the mathematical details, which are routine, may wish to skip to the next section. We capture the requirement for continuous extension by stipulating that a field φ is a continuous function defined between two continua $\varphi : \Omega \rightarrow \Omega'$, where, as usual in topology, a *continuum* is a *nontrivial compact connected set*. Most commonly the domain Ω will be a closed and bounded subset of a Euclidean space, such as a finite interval of the reals, or a disk or rectangle in the plane. The range Ω' will most often be a closed and bounded subset of the real numbers, but complex-valued and vector-valued fields also occur. Finally, it is usually realistic and convenient to restrict our attention to *finite energy* (i.e., $\mathcal{L}_2$) functions, so we can assume fields belong to a Hilbert space.

If φ is a field over Ω and K is a field over $\Omega' \times \Omega$, then we define a kind of product $K\varphi$, which is the continuous analog of a matrix-vector product:

$$K\varphi = \psi, \quad \text{where} \quad \psi(s) = \int_\Omega K(s,t)\varphi(t)\mathrm{d}t.$$

If K is finite energy ($\mathcal{L}_2$) then this product defines an *integral operator of Hilbert-Schmidt type with kernel K*. In MacLennan (subm.-a) we have argued that excitatory synapses and hyperpolarizing inhibitory synapses are *effectively linear*, and so they may be viewed as computing a field product of this kind (Fig. 1A). The field of synaptic efficacies defines the kernel of the operator.

We extend the product notation in the obvious way to more than one argument field. Suppose M is a field over $\Omega \times \Omega_1 \times \cdots \times \Omega_n$ and that φ_k is a

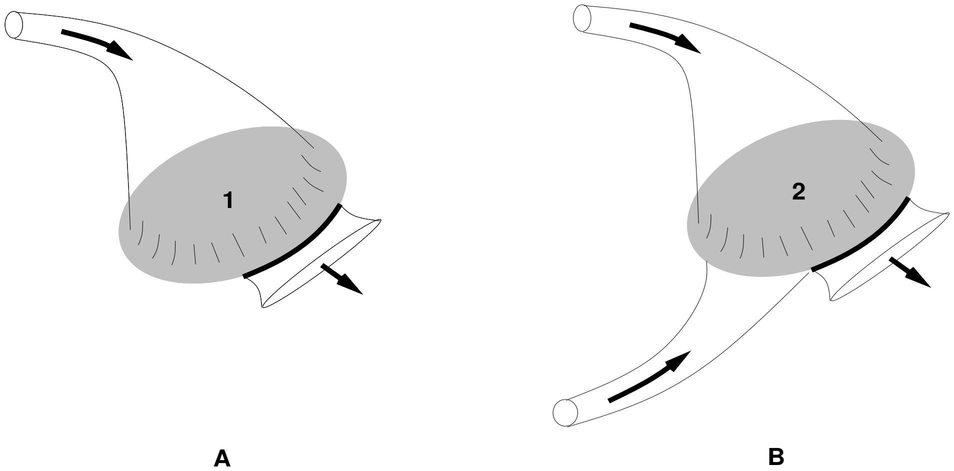

Figure 1: Linear and Bilinear Synaptic Fields. **A.** Excitatory synapses and hyperpolarizing inhibitory synapses can define a linear interaction field L operating on an incoming field φ, $\psi = L\varphi$. More generally these synapses implement an affine field transformation, $\psi = K + L\varphi$. The input and output fields may be defined by spike densities or graded polarizations. (The figure '1' indicates a first-order interaction field.) **B.** Shunting inhibition permits an approximate multiplication between two graded potentials, thus giving a second-order (bilinear) interaction between two fields, $\psi = M\varphi\xi$. More generally such a synaptic field implements a second order interaction, $\psi = K + L_1\varphi + L_2\xi + M\varphi\xi$. The input fields must be graded potentials, but the output field may be represented by graded potentials or spike density. (The figure '2' indicates a second-order interaction field.)

field over Ω_k, $k = 1, \ldots, n$. The product is defined:

$$M\varphi_1\varphi_2\cdots\varphi_n = \psi,$$

where

$$\psi(s) = \int_{\Omega_n} \cdots \int_{\Omega_2} \int_{\Omega_1} M(s, t_n, \ldots, t_2, t_1)\varphi_1(t_1)\varphi_2(t_2)\cdots\varphi_n(t_n)\mathrm{d}t_1\mathrm{d}t_2\cdots\mathrm{d}t_n.$$

The result is a *multilinear operator of Hilbert-Schmidt type with kernel M*; *multilinear* means that it is linear in each of its n arguments.

Notice that the multilinear operator is reduced to the simple product by writing

$$M\varphi_1\varphi_2\cdots\varphi_n = \{\cdots[(M\varphi_1)\varphi_2]\cdots\varphi_n\}.$$

However, all the products after the first, $M\varphi_1$, involve a variable kernel, and so they are in effect bilinear rather than linear operations. This is important, because a shunting inhibition is a second-order operation (it computes a product between two variable quantities), and so a series of second-order operations is of sufficient power to compute any multilinear operator of this kind (Fig. 1B).

There is a well-known theorem in functional analysis that is analogous to Taylor's Theorem in real analysis. It allows expanding a *nonlinear* operator in an infinite series around a fixed field ϖ (MacLennan 1987a; MacLennan 1990):

$$T(\varpi + \varphi) = T(\varpi) + \sum_{k=1}^{\infty} \frac{D_k\varphi^{(k)}}{k!},$$

where

$$D_k\varphi^{(k)} = D_k \underbrace{\varphi\varphi\cdots\varphi}_{k}.$$

The fields D_k are the derivatives of the operator T evaluated at ϖ, $D_k = \mathrm{d}^kT(\varpi)$; these fields give locally-multilinear approximations to T.[7] Although this expansion is "locally good" around ϖ, more often we would like expansions that satisfy global criteria of goodness. Therefore we consider general "polynomial" multilinear expansions of the form:

$$T(\varphi) = K_0 + K_1\varphi + K_2\varphi^{(2)} + K_3\varphi^{(3)} + \cdots.$$

[7]Under the assumptions of field computation, the Fréchet and Gâteaux derivatives are identical (MacLennan 1990).

The kernels K_k are chosen to satisfy or optimize some global criterion, such as minimum error over a training set.

As we've seen, nonlinear operators can be approximated by field "polynomials" of the form $\psi = \sum_k K_k \varphi^{(k)}$, which can be reduced in turn to a sum of first- and second-order interactions. In MacLennan (1987a) we argued that this provides a theoretical basis for universal field computation analogous to the Universal Turing Machine in the theory of discrete computation.[8] Therefore it is especially interesting that first- and second-order interactions can be computed by synaptic fields (Fig. 1), since this suggests that the layers of a neural network may be computing successively higher-order approximations to a nonlinear operator, and series-parallel projections of a field may have a kind of computational universality (Fig. 2).

2.2 Neuronal Field Computation

In this section we'll consider briefly how a number of neural processes can be understood from the standpoint of field computation. For this purpose we will distinguish *neuronal information processing* from *dendritic information processing*. The activity of relatively large cortical areas (i.e., those typically identified anatomically and given names or numbers) can be characterized in terms of *neuronal activity fields*. Most obviously, if $\mathbf{x}$ is the coordinate vector of a neuron in some area Ω, then $D(\mathbf{x})$ could represent the instantaneous spike density of the neuron. Alternately we may view this as a spatiotemporal field (signal or wave) $D(\mathbf{x}, t)$.

Another way of understanding neuronal information processing is in terms of the somatic potential, which represents the integration of a neuron's inputs. Thus we may take $V(\mathbf{x})$ to be the potential of the soma of the neuron with coordinates $\mathbf{x}$. It may be the case that in many areas and in most circumstances, the spike-density field is approximately proportional (up to neuronal precision) to the somatic-potential field, $D(\mathbf{x}) = kV(\mathbf{x})$. This will be the case except for neurons that operate frequently in saturation (i.e., at their minimum or maximum firing rates; see MacLennan subm.-a).

It may be objected that neural cortex is not continuous, but composed of discrete neurons, and of course this is true. But a square centimeter of cortex

[8] In Wolpert & MacLennan (subm.) we show that there is a completely linear field computer that is computationally universal in the sense of Turing.

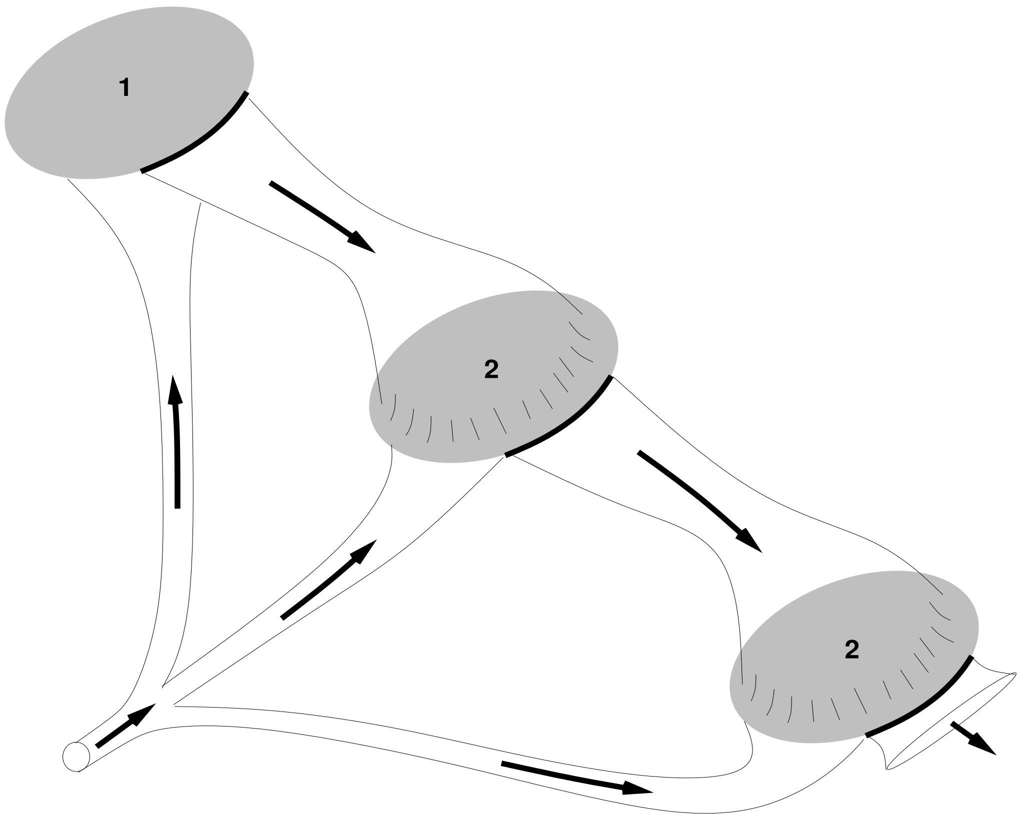

Figure 2: Higher-Order Field Computation in Neural Networks. By projecting in parallel to sequential second-order interaction fields, a neural network may implement a higher-order field polynomial approximating an arbitrary nonlinear operator. Therefore the class of networks of this form exhibit a kind of computational universality. In this case the three-layer network implements the third-degree operator $\psi = K_0 + K_1\varphi + K_2\varphi^{(2)} + K_3\varphi^{(3)}$.

contains approximately 15 million neurons, which is a large enough number to allow the application of continuous mathematics. It is a central tenet of field computation (MacLennan 1987a, 1990) that it does not matter whether the spatial distribution of a quantity is *really* continuous or *really* discrete; to be considered a field it is sufficient that it approximate a continuum well enough to apply continuous mathematics. We believe that for the practical purposes of biological modeling and computer technology, all that matters is whether a phenomenon *looks* continuous or discrete, a methodological tenet called the *Complementarity Principle* (MacLennan subm.-a). We may put it:

> *Continuous models should be practically indistinguishable from approximately-continuous discrete models, and vice versa.*

It has been remarked that neural networks in the brain — as opposed to most PDP models — are neither random nor fully connected (Crick & Asanuma 1986; Pribram 1991, pp. 5–7). Much more common are *neurotopic maps*, topology preserving maps from one cortical area to another.

For an example of neuronal field computation we may take the coordinate transformation that occurs between the retina and its first projection (VI, area 17) in the primary visual cortex.[9] The retinal hemifield is most easily represented in polar coordinates (r, θ), where r represents the radial position from the center of the retina, and θ represents the angle measured clockwise from the horizontal radius of the hemifield (Fig. 3A). Thus $r \in \mathrm{P} = [0, r_{\max}]$, where $r_{\max}$ is the radius of the retina, and $\theta \in \Theta = [-\pi/2, \pi/2]$.

The projection from the retinas to area 17 is *topology preserving* in that regions adjacent on the retina remain adjacent on the cortex. However, there is a *metric distortion* since much more cortical space is devoted to the center of the retina than to the periphery. Indeed, to a first approximation the arrangement of area 17 is as shown in Fig. 3B, which shows logarithmically less cortical distance with increasing distance from the retinal center.[10] If we let (ρ, ϕ) be the coordinates in visual cortex of the point corresponding to

[9] Although this transformation apparently occurs in the projection of the retinal ganglion cells onto the LGN (Berne & Levy 1983, p. 127), we'll take it to be between the retina and VI. Recall that the right visual hemifields of both eyes project to VI in the left hemisphere, and vice versa.

[10] For the time being we are ignoring orientation and velocity sensitivity of cortical cells; they will be considered later.

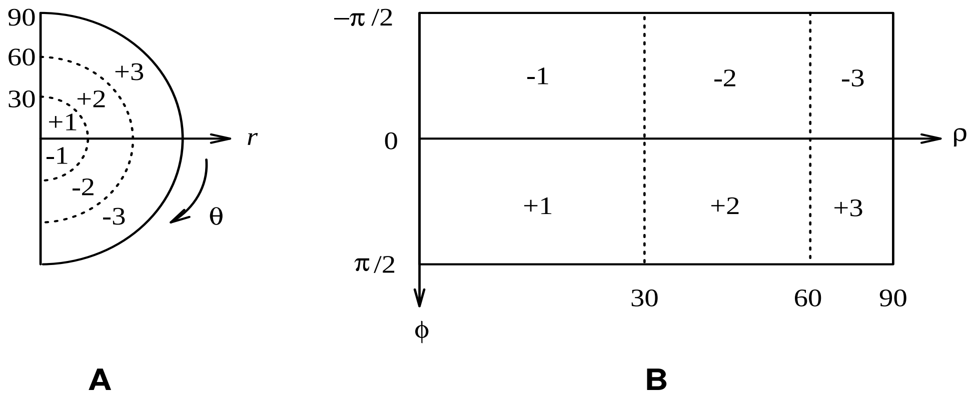

Figure 3: Hemifield Coordinates. **A.** The figure shows the coordinates of the right visual hemifield of either eye. The particular convention chosen for the angles is for convenience in mapping to area 17 (VI) of the visual cortex. **B.** A schematic representation of the retinotopic map of the right hemifields in area 17 (VI) of the left visual cortex. The ρ axis is along the calcarine fissure, from the back of the brain forward (left to right) on the medial surface of the cortex. Radial positions are mapped approximately logarithmically, $\rho = \log r$. Overall the relation between the two coordinate systems is a complex logarithm.

the retinal point at (r, θ), then we see that (ignoring scale factors), $\phi = \theta$ and $\rho = \log r$. Thus the retina-to-cortex map distorts an image φ by applying the transformation:

$$T(\varphi) = \psi, \quad \text{where} \quad \psi(\log r, \theta) = \varphi(r, \theta).$$

Following Baron (1987, pp. 181–186) we note that if we represent retinal position by Cartesian coordinates (x, y), then the coordinate transformation can be expressed

$$\rho = \log \sqrt{x^2 + y^2}, \qquad \phi = \tan^{-1}(y/x).$$

If we express both systems of coordinates by complex numbers, $z = x + iy$, $\zeta = \rho + i\phi$, then this *logmap transformation* is simply a complex logarithm, $\zeta = \log z$. The corresponding field transformation is

$$T(\varphi) = \psi, \quad \text{where} \quad \psi(\zeta) = \varphi(\exp \zeta).$$

The logmap transformation has many information-processing advantages for the visual system (Baron 1987, Ch. 8; Schwartz 1977). In particular, rotations and scale changes of centered retinal images correspond to simple translations of the cortical image.

2.3 Dendritic Field Computation

Shepherd (e.g., 1972, 1978, 1988, 1990a, 1990b) has argued that the synapse, rather than the neuron, should be considered the basic computational element of the brain, and that spatiotemporal relations in the dendritic tree are crucial to understanding synaptic information processing. Further, Pribram (1991, pp. 5–7) has argued that PDP models are a better description of information processing in dendritic nets, which have dense, random interconnection patterns, rather than in neural nets, with their regular topology.

The fields involved in dendritic information processing are predominantly electrochemical fields. Specifically, if $\mathbf{x}$ represents the location of an active site in the dendritic arbor, such as a synapse, then $\varphi(\mathbf{x})$ most commonly represents the membrane potential, but it could also represent the concentration of a chemical species, such as an ion or neurotransmitter. Of course, there has been much work, from Hodgkin and Huxley's day, that models the detailed dynamics of these processes, but that is not our concern here. Since we are interested in the general structure of information processing in the brain, it is sufficient that such fields exists; we would like to understand their possible role in information representation and processing.

The electrochemical dynamics of nervous tissue is without doubt complex, so the challenge is to find simplifying approximations that capture the essence of information processing and avoid irrelevant detail. To this end we have been investigating a *linear system* model of dendritic information processing (MacLennan subm.-a; MacLennan & Pribram in prep.). The mathematical advantages of a linear model are obvious, but do we have any basis for assuming it? We believe that the evidence supports linear models of both excitatory synapses and hyperpolarizing inhibitory synapses, and bilinear models of shunting inhibition (MacLennan subm.-a; MacLennan & Pribram in prep.). Although there is widespread opinion that linear systems are computationally impotent (e.g., Poggio & Torre 1981; Reichardt & Poggio 1981), we show below that dynamic linear systems can accomplish significant information processing (cf. also Wolpert & MacLennan subm.).

We also consider the present theory a simplification of the neural wave equation developed by Pribram, Yasue and Jibu (Pribram 1991, Apps. A–G), which is also a linear model. Specifically, by assuming a discrete set of interaction sites, the model becomes a *lumped-parameter* system, which means that its dynamics can be described by ordinary differential equations rather than partial differential equations. A significant simplification results from this assumption, which is justified by our Complementarity Principle, and we anticipate that this simplification will help us to go beyond the dynamics of dendritic interactions, and to understand their *function*.

Let ψ be some time-varying field relevant to dendritic information processing (membrane polarization at synapses would be an example). Based on the linear-system assumption, we take its dynamics to be defined by a nth-order integro-differential equation with kernels F_k:

$$\psi^{(n)} = \sum_{k=0}^{n-1} F_k \psi^{(k)} + \text{input drive}.$$

As usual this equation can be reduced to a system of first-order equations by introducing additional field variables ψ_k corresponding to the time-derivatives of ψ:

$$\begin{aligned} \psi &= \psi_0, \\ \dot{\psi}_k &= \psi_{k+1}, \quad k = 0, \ldots, n-1, \\ \dot{\psi}_{n-1} &= \sum_{k=0}^{n-1} F_k \psi_k + \text{input drive}. \end{aligned}$$

If these state variables are combined into a field-vector $\boldsymbol{\psi} = (\psi_0, \ldots, \psi_{n-1})$ and the kernels $F_0, \ldots, F_{n-1}$ into a kernel-vector F, then the system can be described by a single field-vector differential equation:

$$\dot{\boldsymbol{\psi}} = F\boldsymbol{\psi} + \text{input drive}.$$

To complete our description of the linear system, we must describe how it is driven by the input φ and how it in turn drives the output ω. Then we have a linear system of the form (Fig. 4):

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

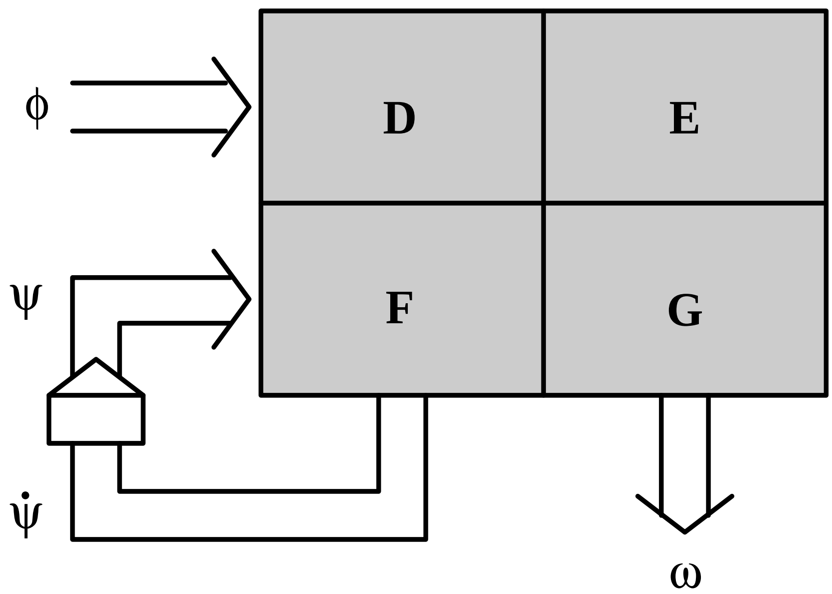

Figure 4: The Dendritic Net as a Dynamic Linear System. The input field φ drives both the state fields $\boldsymbol{\psi}$ through kernel D and the output field ω through kernel E. Kernel F determines the feedback among the state fields, and kernel G governs the contribution of the state fields to the output field. The house-shaped figure represents a field integration (with respect to time) of the state fields.

Such a system will exhibit *resonances*, the number of which is on the order of the size of the state fields, that is on order of the number of interaction sites. Since there may be 5 000 to 200 000 synapses in the dendritic arbor of a single neuron, it's not implausible to assume that such a dendritic net may have thousands of resonances. We consider elsewhere some kinds of information processing that such a system can implement, including the self-organization of hierarchically-structured spatiotemporal matched filters (MacLennan subm.-a).

Although our emphasis here is on the linear processes, it will be worthwhile to say a little about the functional role of nonlinearities in dendritic information processing. We have already mentioned the bilinearity of shunting inhibition; one possible function it could serve (aside from a simple and-not gate) is to implement bilinear operators such as convolutions and correlations over both space and time, whose information-processing potential is manifest. Also, as explained previously, bilinear operations are sufficient for an important kind of computational universality. One of the best known nonlinearities in the behavior of neurons is the generation of action potentials. According to Shepherd (1988, p. 137), an action potential causes an *antidromic* electrical signal, which is transferred efficiently into the dendritic spines. We have suggested (MacLennan subm.-a) several possible roles for this signal, including the (1) triggering or enhancing of Hebbian learning, (2) the top-down "priming" of dendrites (i.e., the creation of top-down expectations), and (3) pattern completion. We've also shown how the antidromic electrotonic flow can lead to the self-organization of recursive spectral-density matched-filters, which have many possible applications in neural information processing.

Hameroff (1987) has suggested that Boolean and automata-like processes could occur in the cytoskeletons of neurons, and that this could be "where the action is" so far as information processing is concerned. We suggest that field computation may be a better model of cytoskeletal processing — if it exists — since the large number of elements (~1625/micron) makes the microtubule a good approximation to a field.[11] Local values of the field could be represented, for example, by rate of conformational change of the

[11]Microtubules are composed of spirals of tubulin dimers, 13 around the circumference, which are 8 nm long in the direction of the axis (Hameroff 1987, p. 106). Therefore, a one-micron length of microtubule contains 1000nm / 8nm spirals, each containing 13 dimers; $13 \times 1000/8 = 1625$.

microtubule-associated proteins. The question remains, of course, whether the dynamics are approximately linear, but if they are, then much of the theory presented here would apply unchanged.

3 Gabor-like Representations

In this section we will consider some Gabor-like field representations that may be important in sensory and motor systems.

3.1 Vision

3.1.1 Spatial Gabor Wavelets

There is now considerable evidence that the receptive fields of simple cells in the primary visual cortex correspond (up to synaptic precision) to the even- or odd-symmetric parts of two-dimensional Gabor functions $\gamma_{\mathbf{q}}$, which suggests that Gabor functions are the representational primitives of the primary visual cortex.[12] (See MacLennan 1991 for a review.)

The problem is that the Gabor functions are not orthogonal, so the coefficients $c_{\mathbf{q}}$ of a Gabor expansion of a field φ,

$$\varphi = \sum_{\mathbf{q}} c_{\mathbf{q}} \gamma_{\mathbf{q}},$$

cannot be computed by a simple inner product, $c_{\mathbf{q}} \neq \langle \varphi, \gamma_{\mathbf{q}} \rangle$. Nevertheless, inner products with Gabor functions are what the simple cells seem to compute.[13]

[12]An N-dimensional Gabor function is a Gaussian-modulated complex-exponential,

$$\gamma_{S\mathbf{pu}}(\mathbf{x}) = \exp\{-\pi \| S(\mathbf{x} - \mathbf{p}) \|^2\} \exp[2\pi i \mathbf{u} \cdot (\mathbf{x} - \mathbf{p})].$$

All the vectors are N-dimensional. The parameter $\mathbf{p}$ determines the function's location in N-space; the wave vector $\mathbf{u}$ determines its modulatory frequencies, or position in N-dimensional spectral space, and the orientation of the function in N-space. The parameter S is a diagonal matrix which defines the function's *aspect ratio*, or shape in N-space. When it is not necessary to distinguish the parameters, we simply write $\gamma_{\mathbf{q}}(\mathbf{x})$. The Gabor functions form an *anisotropic (oriented) wavelet family.*

[13]This is an oversimplification, as noted in MacLennan (1991, n. 23). The input to area 17 is from the retina via the LGN, which have already represented the image in terms of

Although there are theoretical reasons to expect representation in terms of Gabor functions (they are optimal in terms of the Gabor Uncertainty Principle), we must keep in mind that they are mathematical objects and cannot be instantiated perfectly in the biology. Therefore, objections against the Gabor functions, such as that they have noncompact support, are not relevant in a biological context. It's true that the Gaussian envelope extends to infinity, but 99.7% of its area is within three standard deviations of its mean, and 99.994% within four. Thus the Gabor functions are practically indistinguishable from functions with compact support. The conclusion we draw is that the theory is underdetermined by the biology, and so we can choose to model the receptive fields by Gabor functions, if it is mathematically convenient to do so.

Another problem with the Gabor functions is that they are complex-valued, and therefore not representable by real-valued membrane potential, spike densities, etc. Although there is evidence (Pollen & Ronner 1981) that simple cells occur in conjugate pairs with receptive fields representing the (real valued) odd- and even-symmetric parts of the Gabor function, Stork & Wilson (1990) have objected that these real functions do not minimize the Gabor uncertainty, and therefore that they should not be given special status. The real-valued functions that achieve the minimum are the Hermite functions, as shown by Gabor and proved more carefully by Stork & Wilson. Nevertheless, the Hermite functions are sufficiently like the real parts of the Gabor functions, that even this difference may be insignificant.

Although Gabor functions are nonorthogonal and so cannot be a basis, under biologically plausible conditions they do form a *frame*, for which the inner products fulfill a similar role to that for bases (MacLennan 1991). Thus one possibility is that the higher levels of the visual system simply operate in terms of a nonorthogonal representation. Just because orthogonality is mathematically convenient doesn't imply that it's biologically convenient. As Daugman (1988) observed, nonorthogonality is ubiquitous in

radial basis functions (center-surround receptive fields). Therefore, the Gabor receptive fields observed in striate cortex reflect the combined effect of retina, LGN and area 17 on the image.

sensory and motor systems, which we should expect from the *Robustness Principle* (MacLennan subm.-a):

> *No biological process can depend on an absolute mathematical property.*

This tells us that biological processes cannot depend on exact orthogonality. In fact, approximate orthogonality is in many ways preferable to exact orthogonality (Kainen submitted; MacLennan subm.-a). Further, given the imprecision of biological computation, the correction of coefficients by relaxation (such as described below) can be expected even for formally orthogonal representational primitives.

In MacLennan (1991) we argued that such a relaxation process was unlikely in neural networks in which signaling is mediated by impulses, since the rate of information transmission is too slow, but that it was feasible in local circuits in dendritic nets, where signaling may be mediated by graded potentials. In MacLennan (subm.-a) we showed that dendritic iteration can be expected to proceed 20 to 200 times as fast as neuronal iteration, and furthermore that dendritic nets have the computational power to implement a linear system that relaxes to the coefficients of a nonorthogonal representation, such as the Gabor (see also MacLennan & Pribram in prep.). In this linear system, the driving matrix D is proportional to the array of elementary fields $\varrho_{\mathbf{q}}$,

$$D_{\mathbf{q}} = \eta \varrho_{\mathbf{q}},$$

for example even- or odd-symmetric Gabor functions, which correspond with the observed receptive field profiles. The feedback matrix or interaction field F is proportional to all the inner products between the elementary fields (i.e., the Gram matrix),

$$F_{\mathbf{qr}} = -\eta \langle \varrho_{\mathbf{q}}, \varrho_{\mathbf{r}} \rangle,$$

and has a decorrelating effect which takes care of the nonorthogonality.

Pattison (1992) has independently proposed the same relaxation algorithm for Gabor or other nonorthogonal representations, but claims that it could be implemented in neural networks. In particular he assumes that the coefficients are represented by instantaneous firing frequencies, but neglects the finite interval — given by the Gabor Uncertainty Principle — required to represent a coefficient to a given precision (MacLennan 1991, subm.-a;

MacLennan & Pribram in prep.). In fact, approximately 10^k msec. are required for k digits of precision. Nevertheless he estimates that the relaxation process will require at least 50 msec. per iteration, which he concedes is inconsistent with simple cell response times observed by Jones & Palmer (1987). We suggest that the inconsistency is eliminated by assuming that the relaxation takes place through graded interactions in the dendritic net, which have a delay of about 1 msec.

We consider briefly the representation of the Gabor coefficients in visual cortex. We've already seen that the logmap transformation converts retinal coordinates (r, θ) into VI coordinates (ρ, ϕ); recall Fig. 3. If the VI field were simply a representation of light intensity at the retina, then it would be a scalar field over the two-dimensional domain $\mathrm{P}\times\Theta$, but we've seen that it's a scalar field over three dimensions: two for retinal location and one for Gabor-field orientation.[14] The possible orientations of the receptive fields range from 0 to 2π so we define $\Phi = [0, 2\pi]$; it is not simply $[0, \pi]$ since the odd-symmetric receptive fields are asymmetric across their edge. Therefore, VI must represent a scalar field defined over three-dimensions, $\mathrm{P}\times\Theta\times\Phi$. How can this be represented in the essentially two-dimensional cortex?

This is a standard problem in field computation (MacLennan 1990), and there are a number of ways of mapping a higher dimensional field into a lower dimensional space. One is to simply cut up the field along one or more of the dimensions, and arrange the resulting lower-dimensional fields next to one another. This is exactly what we find in the orientation columns of striate cortex. The orientations $[0, 2\pi]$ run through an entire cycle in a space of about 2mm and repeats thereafter (Baron 1987, p. 153). Therefore, wherever the mathematical function requires a field of dimension greater than two, we expect to see a striate or columnar structure in the cortex.

This arrangement has another useful effect in that it leads to texture (oriented spatial frequency) being represented is a similar way to color: the orientation columns are arranged similarly to the *cortical pegs* or *blobs* that respond to color (Shepherd 1988, pp. 348–349). This is reasonable since both texture and color are "extended" properties.

[14]There is actually a fourth dimension, representing spatial frequency and receptive-field size, which are closely correlated. For simplicity this will be ignored.

3.1.2 Spatiotemporal Gabor Wavelets

The research which led to the Gabor representation (Gabor 1946) was motivated in part by the observation that our perception of sound is simultaneously of duration and pitch. Thus the Gabor representation captures the local temporal structure of the sound. The preceding discussion of vision can be interpreted to mean that our visual perception is simultaneously of extension and texture (oriented spatial frequency). Thus the Gabor representation captures the local spatial structure of a scene. On the other hand, we also have an immediate visual awareness of motion, and we know that some cells in the visual cortex respond to motion, so we might wonder if our visual system uses a Gabor representation in both the spatial and temporal domains.

In MacLennan (1991) we suggest three-dimensional Gabor functions, with two space dimensions and one time dimension, as possible representational primitives in vision. In addition to orientation in space, these functions may be oriented in space-time, which gives them receptive fields responsive to textural motion localized in space and time (Fig. 5). Temporal localization suggests that visual system operates cyclically, which is compatible with the use of relaxation to calculate Gabor coefficients. The idea is that the basic cycle is (1) acquire the image; (2) calculate its Gabor coefficients by relaxation; (3) forward the coefficients on to the next stage of processing. It's possible that the frequency of this cycle is the alpha rhythm, which is the principal rhythm of the occipital cortex. It might seem that representational primitives oriented in space-time are of only theoretical interest, but in fact they have a simple neural implementation in terms of spatial Gabor functions, analogous to the well-known construction of moving-edge detectors from static-edge detectors. Although spatiotemporal Gabor functions seem to be consistent with neurophysiological data, more research is needed to establish their presence.

3.2 Other Sensory and Motor Systems

The ability of Gabor and similar locally-Fourier representations to capture temporal structure suggests that we look for them in other sensory and motor systems; perhaps they are a general representational principle in the brain. For example, observe that an auditory signal, from the cochlea on, is repre-

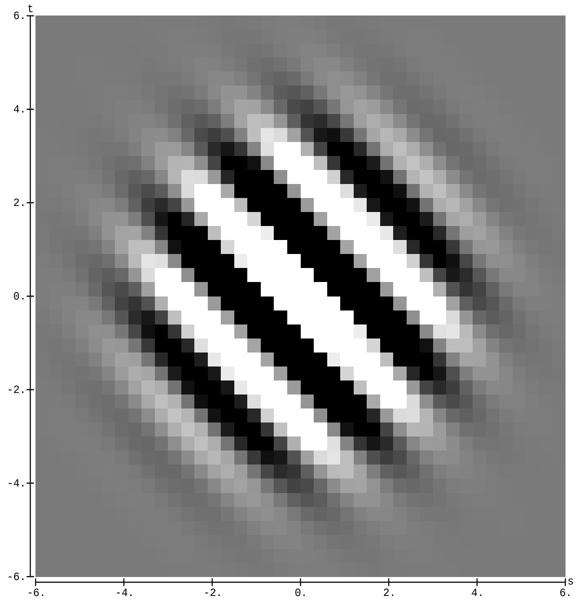

Figure 5: Depiction of Spatiotemporal Gabor Function. The figure shows a slice through the even (cosine) part of a 3D Gabor function. The vertical axis represents time and the horizontal axis is along the spatial orientation of the function. This filter is selective for fringes of a specific frequency moving at a specific velocity in a localized region of space and time. This particular filter responds maximally to fringes with a frequency $1/\sqrt{8}$ moving with a velocity of 1 to the left (arbitrary units).

sented by a spatiotemporal wave, in which instantaneous frequency — the local Fourier transform — is mapped to spatial location. Therefore, 2D spatiotemporal Gabor functions will capture temporal changes of intensity and pitch — what we might call the rhythm and melody of the sound.

Temporal structure is also very important in motor activity, and so we will consider the possibility of Gabor-like representations there. It is reasonable to treat as a field the activity of a large system of motor neurons or of a region of motor cortex. Motor control is then accomplished by generating an appropriate spatiotemporal wave (in a feedback loop, of course). Since such a signal may be represented as a linear superposition of Gabor-like functions, we will investigate the effect of such a representation of motor signals.

Considered as a *generative field* rather than a receptive field, a Gabor-like function amounts to an amplitude-controlled, time-bounded rhythm generator (Fig. 5). More concretely, a Gabor function $\gamma_{\mathbf{q}}$ generates, during a given time interval Δt, a burst of waves of a given frequency and direction across a field of motoneurons (such as a somatotopically mapped region). Higher levels of motor control have the task of generating the Gabor coefficients. More global control mechanisms can adjust the size of the Gabor functions, which affects both their frequency and their spatial and temporal extent.

4 Cognition

In this section we touch on the role field computation may play in higher cognitive function; our focus will be on *intentionality*.

4.1 Intentions

4.1.1 History

The notions of *intentionality* and *intention* are central to the modern philosophy of the mind, so it will be useful to review the meanings of these terms before considering *intentional fields*.

The basic meaning of Latin *intendo* is to stretch toward, point at, or to direct one's mind toward, and at least from the time of Cicero the related noun *intentio* could refer to acts of stretching, reaching or concentrating one's attention (*Oxford Latin Dict.*, s.vv.). The basic idea is an active process of directing the mental faculties.

The Medieval Schoolmen chose *intentio* to translate Arabic *ma'nā* (a meaning, thought, signification or notion) in the works of Avicenna. Later, Ockham defined an intention as "something in the soul capable of signifying something else," or more briefly, a sign in the mental discourse (*Summa Logicae* I § 12). Although this definition is limited by Ockham's linguistic view of cognition, it captures the idea that an intention is a mental representation that refers to something outside of itself (either in the world or in another mental representation).

Brentano (in his *Psychology from an Empirical Standpoint*) resurrected the medieval notion of intention, and used it to refer to the ability of consciousness to refer out of itself and be directed toward something, that is, the characteristic of consciousness that it is consciousness *of* something. Husserl borrowed the term from Brentano and used it with this sense. In accord with its methodological biases, Anglo-American philosophy gave intentionality a linguistic interpretation. An intentional proposition is one which has another proposition as its content, and so it expresses a certain "attitude" toward that proposition.

All these definitions have in common the idea of *selecting* or "foregrounding" some aspects of a mental representation with respect to others, which are left in the background. A "mental sign" picks out some aspect of the external or internal world (in Scholastic terms, a first or second intention); in Brentano and Husserl's terms, consciousness is directed at some aspects of experience; in linguistic terms, an intentional expression has a particular content. The essence of selection is a decrease in entropy, for by making some things more likely to be processed relative to others, we shift the probability distribution away from the uniform distribution, which has maximum entropy, to a nonuniform, lower entropy distribution (MacLennan 1988, pp. 172–173). Thus an intention organizes a representation with respect to an intended functional role. Intentions, as understood here, are closely related to the focus of attention (Pribram 1991, pp. 119–120, 219).

4.1.2 Field Representation

Next we will consider a possible theory of intentions in terms of field computation. We define an *intentional field* to be a $[0, 1]$-valued field over any domain Ω. If φ is any field over the same domain Ω as an intentional field ν, and $F(\nu, \varphi)$ is an operation on the pair of fields (ν, φ), then ν is called an *intention of type F toward* φ.[15] This definition reflects the fact that an essential part of an intention is its *function* or *end*; it is an intention to treat something in a particular way (e.g., to notice it, to be surprised by it, to fear it, to avoid it, to orient toward it, to seize it). Thus the operator F represents the function of the intention, which I call its *kind*. (In biological terms, F might correspond to a specific brain area and ν to a field over that area.) Each intention also has a particular *content*, which is the region of the image towards which the operation is directed; the content is given by the pair of fields (ν, φ). The intention field ν can be interpreted as a probability distribution selecting certain regions of φ for more-likely processing by F. In this sense an intention functions like a continuous analog of a programming-language pointer.

Some examples may make this idea clearer. Visual intentions are characteristic responses or attitudes to the content of the visual field. Thus we may be surprised, either by the presence of some object or by its absence. Also, we may react to perceived objects with fear or with comfortable familiarity. These are not purely visual intentions, since they typically involve nonperceptual evaluations. Purer examples include the tracking of a moving object, or the sudden focus of our attention on an unexpected movement. It is clear in both of these how the relevant intention field could be computed from relatively low level perceptual fields.[16] Auditory intentions are similar, including, for example, the ability to focus on a particular sound, such as a voice, among many other sounds.

For an example of a nonperceptual intention, consider our awareness of an object not visible to us (e.g. behind our back, in a closet). We suppose that spatial awareness is represented in several frameworks, such as the egocentric, centered on the body, and the allocentric, centered elsewhere (Bryant 1990; Bryant et al. 1992; Franklin & Tversky 1990), and that these representations are not unlike abstract sensorimotor representations. Spatial intentions then refer to locations within these spaces, and establish functional relations with

[15] I will generally use ν (from $\nu o \acute{\epsilon} \omega$ = intend, discern, notice, remark) for intentional fields. Similarly I will generally use φ (from $\varphi \acute{\alpha} \nu \tau \alpha \sigma \mu \alpha$ = mental image) for fields representing images of any other sort.

[16] For example, the spatiotemporal Gabor representation makes it easy to detect sudden changes in the motion in a local area. We simply apply a low pass filter to eliminate slow changes. Of course that's not the whole story, since animals also habituate to rapid changes.

those locations (e.g., intent to move an arm to that place, noting presence or absence of an object in that place). Spatial intentions may be translated from one reference frame to another in the same way as images.

Orienting reactions can be understood in this context. An unexpected sight, sound or touch creates a perceptual intention (ν, φ) in a functional area F. The intentional field ν is translated from sensoricentric to egocentric coordinates in the same way as a perceptual image. This "surprise" intention has the functional role of being translated into an intention to move, which can be used to compute a spatiotempral motor field to control the motion.

4.2 Abstract Reason

Johnson-Laird and Byrne have argued that abstract reason is accomplished by manipulating mental models rather than formal symbols (Johnson-Laird & Byrne 1991, in press). We suggest that these models are just abstract, multimodal sensorimotor images, including intentions (MacLennan subm.-b), and we anticipate that manipulation of mental models can be described in terms of field computation. To give a very rough idea, we consider one of the formal logic problems studied by Johnson-Laird & Byrne:

> There is a cross if and only if there is a circle. There is not a cross. What follows?

Let φ be an (abstract) image of both the cross and the circle. Let ν_{xp} be an intentional field indicating that the cross is present, and ν_{op} that the circle is present. Then $\nu_{\mathrm{p}} = \nu_{\mathrm{xp}} + \nu_{\mathrm{op}}$ is an intention referring to both and indicating their presence.[17] To show the absence of the objects, we postulate intentional fields ν_{xa} and ν_{oa}, and their sum $\nu_{\mathrm{a}} = \nu_{\mathrm{xa}} + \nu_{\mathrm{oa}}$.

The first premiss (the biconditional) results in the construction of two models, $(\nu_{\mathrm{p}}, \varphi)$ and $(\nu_{\mathrm{a}}, \varphi)$, which are held in working memory (as fields in prefrontal cortex).

The second premiss, "there is not a circle," results in the model $(\nu_{\mathrm{oa}}, \varphi)$. In an attempt to merge this with each of the previous models, it's found to be

[17]In fact, these intentions mean that these are the only allowed models in which these things are present; thus they include Johnson-Laird & Byrne's "exhaustive representation tag"; there are other intentions, not discussed here, that do not bear that tag (i.e., they have a different functional role).

incoherent with the first, since the intentions ν_{oa} and ν_{p} treat φ inconsistently (as can be seen by noting $\nu_{\mathrm{oa}}\nu_{\mathrm{p}} \neq 0$). The only other model, that the cross and circle are both absent, is consistent with the second premiss, so it is the conclusion (or, more accurately, the model that the conclusion expresses).

5 Conclusions

We have argued that the limitations of the traditional representation of knowledge as discrete symbols can be avoided by a reformulation in terms of continuous images. We proposed the *simulacrum* as a topological model of continuous knowledge representation and processing, and discussed *field computation* as a specific instance of it. We showed that both neuronal and dendritic information processing can be understood in the context of field computation. In the course of this we considered the possible role of spatiotemporal Gabor wavelet representations in sensorimotor systems, and suggested a field representation of intentions. Overall we hope that this paper has shown the potential contribution of continuous computation — and especially field computation — toward understanding the mind and brain.

6 References

Baron, R. J. (1987). *The cerebral computer: An introduction to the computational structure of the human brain.* Hillsdale NJ: Lawrence Erlbaum.

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

Blum, L. (1989). *Lectures on a theory of computation and complexity over the reals (or an arbitrary ring)* (Report No. TR-89-065). Berkeley CA: International Computer Science Institute.

Blum, L., Shub, M., & Smale, S. (1988). On a theory of computation and complexity over the real numbers: NP completeness, recursive functions and universal machines. *The Bulletin of the American Mathematical Society*, **21**: 1–46.

Bryant, D. J. (1992). A spatial representation system in humans. *PSYCOLOQUY*, **3**(16) `space.1`.

Bryant, D. J., Tversky, B., & Franklin, N. (1992). Internal and external spatial frameworks for representing described scenes. *Journal of Memory and Language* **31**: 74–98.

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* (pp. 333–371). Cambridge MA: MIT Press.

Daugman, J. G. An information-theoretic view of analog representation in striate cortex. In: E. L. Schwartz (Ed.), *Computational Neuroscience* (pp. 403–423). Cambridge MA: MIT Press.

Dennett, D. (1987). *The intentional stance.* Cambridge MA: MIT Press.

Dennett, D. (1988). The intentional stance. *Behavioral and Brain Sciences*, **11**: 495–546.

Franklin, S., & Garzon, M. (1990). Neural computability. In: O. M. Omidvar (Ed.), *Progress in neural networks* (Vol. 1, pp. 127–145). Norwood NJ: Ablex.

Franklin, N., & Tversky, B. (1990). Searching imagined environments. *Journal of Experimental Psychology: General*, **119**: 63–76.

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

Garzon, M., & Franklin, S. (1989). Neural computability II (extended abstract). In: *Proceedings, IJCNN International Joint Conference on Neural Networks* (Vol. 1, pp. 631–637). New York NY: Institute of Electrical and Electronic Engineers.

Garzon, M., & Franklin, S. (1990). Computation on graphs. In: O. M. Omidvar (Ed.), *Progress in neural networks* (Vol. 2, Ch. 13). Norwood NJ: Ablex.

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

Harnad, S. (1990). The symbol grounding problem. *Physica D*, **42**: 335–346.

Harnad, S. (in press). Grounding symbols in the analog world with neural nets: A hybrid model. *Think.*

Hartley, R., & Szu, H. (1987). A comparison of the computational power of neural network models. In: M. Caudill & C. Butler (Eds.), *Proceedings, IEEE First International Conference on Neural Networks* (Vol. 3, pp. 17–22). New York NY: Institute of Electrical and Electronic Engineers.

Johnson- Laird, P. N., & Byrne, R. M. J. (1991). *Deduction.* Hillsdale NJ: Lawrence Erlbaum.

Johnson-Laird, P. N., & Byrne, R. M. J. (in press). Precis of *Deduction. Behavioral and Brain Sciences*, in press.

Jones, J. P., & Palmer, L. A. (1987). The two-dimensional spatial structure of simple receptive fields in cat striate cortex. *Journal of Neurophysiology*, **58**: 1187–1211.

Kainen, P. C. (submitted). Orthogonal dimension and tolerance.

Köhler, W. (1940). *Dynamics in Psychology.* New York NY: Liveright.

MacLennan, B. J. (1987a). 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. (1987b). *Field computation and nonpropositional knowledge* (Report NPS52-87-40). Monterey CA: Naval Postgraduate School, Computer Science Department.

MacLennan, B. J. (1988). Logic for the new AI. In: J. H. Fetzer (Ed.), *Aspects of Artificial Intelligence* (pp. 163–192). Dordrecht NL: Kluwer Academic Publishers.

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. (1992a). Synthetic ethology: An approach to the study of communication. In: C. G. Langton, C. Taylor, J. D. Farmer and S. Rasmussen (Eds.), *Artificial Life II* (pp. 631–658). Redwood City, CA: Addison-Wesley.

MacLennan, B. J. (1992b). *Research Issues in Flexible Computing: Two Presentations in Japan* (Report No. CS-92-172). Knoxville TN: University of Tennessee, Knoxville, Computer Science Department.

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. (submitted-a). Information processing in the dendritic net.

MacLennan, B. J. (submitted-b). Visualizing the possibilities (review of Johnson-Laird & Byrne 1991).

MacLennan, B. J. (submitted-c). Grounding analog computers (commentary on Harnad in press).

MacLennan, B. J., & Burghardt, G. M. (submitted). Synthetic ethology and the evolution of cooperative communication.

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

Nemytskii, V. V., & Stepanov, V. V. (1989). *Qualitative theory of differential equations*. New York NY: Dover; also (1960) Princeton NJ: Princeton University Press.

McCulloch, W.S., and Pitts, W. (1943). A logical calculus of the ideas immanent in nervous activity. *Bulletin of Mathematical Biophysics*, **5**: 115–133.

Pattison, T. R. (1992). Relaxation network for Gabor image decomposition. *Biological Cybernetics*, **67**: 97–102.

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.

Pollack, J. B. (1987). *On connectionist models of natural language processing* (Ph.D. dissertation). Urbana IL: University of Illinois; also report MCCS-87-100, Las Cruces NM: New Mexico State University, Computing Research Laboratory.

Pollen, D., & Ronner, S. (1981). Phase relationships between adjacent simple cells in the visual cortex. *Science*, **212**: 1409–1411.

Pour-El, M. B., & Richards, I. (1979). A computable ordinary differential equation which possesses no computable solution. *Annals of Mathematical Logic*, **17**: 61–90.

Pour-El, M. B., & Richards, I. (1981). The wave equation with computable initial data such that its unique solution is not computable. *Advances in Mathematics*, **39**, 215–239.

Pour-El, M. B., & Richards, I. (1982). Noncomputability in models of physical phenomena. *International Journal of Theoretical Physics*, **21**: 553–555.

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

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

Schwartz, E. L. (1977). Spatial mapping in the primate sensory projection: Analytic structures and relevance to perception. *Biological Cybernetics*, **25**: 181–194.

Scott, D. S. (1970). Outline of a mathematical theory of computation. In: *Proceedings of the Fourth Annual Princeton Conference on Information Sciences and Systems* (pp. 160–176). Princeton NJ: Princeton University, Department of Electrical Engineering.

Scott, D. S. (1971). The lattice of flow diagrams. In: E. Engeler (Ed.), *Symposium on the Semantics of Algorithmic Languages* (pp. 311–366). Berlin: Springer-Verlag.

Scott, D. S. (1973). Lattice-theoretic models for various type-free calculi. In: P. Suppes, L. Henkin, A. Joja & G. C. Moisil (Eds.), *Logic, Methodology and Philosophy of Science IV* (pp. 157–187). Amsterdam NL: North-Holland.

Scott, D. S., & Strachey, C. (1971). Towards a mathematical semantics for computer languages. In: J. Fox (Ed.), *Proceedings of the Symposium on Computers and Automata* (pp. 19–46). New York NY: Polytechnic Institute of Brooklyn Press.

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.). (1990a). *The synaptic organization of the brain*, third edition. New York NY: Oxford University Press.

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

Stannett, Mike (1990). X-machines and the halting problem: Building a super-Turing machine. *Formal Aspects of Computing*, **2**: 331–341.

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**: 1362–1373.

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