

Ballard
980

TECH REPORT

780
11467

**Cortical Connections and Parallel
Processing: Structure and Function**

Dana H. Ballard
Computer Science Department
University of Rochester

TR 133 (revised)
January 1985

Abstract

The cerebral cortex is a rich and diverse structure that is the basis of intelligent behavior. One of the deepest mysteries of the function of cortex is that neural processing times are only about one hundred times faster than the fastest response times for complex behavior. At the very least, this would seem to indicate that the cortex does massive amounts of parallel computation.

This paper explores the hypothesis that an important part of the cortex can be modeled as a connectionist computer that is especially suited for parallel problem solving. The connectionist computer uses a special representation, termed value unit encoding, that represents small subsets of parameters in a way that allows parallel access to many different parameter values. This computer can be thought of as computing hierarchies of sensory-motor invariants. The neural substrate can be interpreted as a commitment to data structures and algorithms that compute invariants fast enough to explain the behavioral response times. A detailed consideration of this model has several implications for the underlying anatomy and physiology.

CARLSON LIBRARY

Table of Contents

1. Introduction
2. Overview
 - 2.1 Representing Possible Invariants
 - 2.2 Computing Specific Invariants
3. A Review of Cortical Organization
 - 3.1 Cortical Columns: A Uniform Processing Architecture
 - 3.2 Functionally Different Cortical Areas
 - 3.3 Abstraction Hierarchies
 - 3.4 Spatio-Temporal Channels
4. Value Unit Constraints on Cortical Organization
 - 4.1 Relating Function to Cortical Topography in Striate Cortex
 - 4.2 Primary and Secondary Indices
 - 4.3 Space Limitations Require Different Functional Areas
 - 4.4 Effects of Changing Indexing Schemes on Anatomical Connections
 - 4.5 Developmental Constraints on Cortical Structure
5. Computing with Cortico-Cortical Connections
 - 5.1 Shape Perception
 - 5.2 Motion Perception
6. Associating Value Units in Different Modalities
7. Conclusion
8. Acknowledgements
9. References

1. Introduction

Tremendous progress could be realized in the neurosciences by the introduction of information processing models that would relate biological and behavioral data. Information processing models (IPMs) have been successful in describing behavior at levels of abstraction useful to psychologists and workers in artificial intelligence, but have been much less successful in producing models that are directly useful to neuroscientists. One reason for this is that most IPMs are based on a conventional computer and animal brains do not compute like a conventional computer. In the animal brain, comparatively slow (millisecond) neural computing elements with complex, parallel connections form a structure which is dramatically different from a high-speed (nanosecond), predominantly serial machine. Much of current research in the neurosciences is concerned with tracing out these connections and with discovering single unit responses to complex stimuli. However, a crucial next step is to characterize neural function at the more abstract level that is related to behavior.

To characterize our endeavor it might help to consider a complementary approach: temporarily side-step the biological issues and study the *abstract computational problems* that must be solved [Marr and Poggio, 1976; Brady, 1982; Marr, 1982]. The objective of this tack is to find useful abstract descriptions of the computations that are being performed without reducing them to anatomy. The level of formulation is in terms of symbolic constraints and algorithms for solving them. Subsequently, this class of models must be complemented by a description of how the brain *implements* the solution to specific computational problems. This paper is aimed at this descriptive level. The implementation approach may be thought of as one of synthesis, and is logically at the boundary between biology and computer science. Neurobiological models describe the anatomy and physiology accurately while postponing the problem of computation. Computational models stress the abstract nature of the problems that must be solved, postponing the problem of implementation. In the implementational approach, one must choose a description of brain architecture that both describes how the abstract problems can be solved and is neurobiologically plausible.

One school of thought, termed *connectionist*, holds that the essential components of the abstract level can be described in terms of the synaptic contacts of networks of neurons. In other words, the functionality can be directly and usefully related to neural interconnection patterns. Early connectionist models [McCulloch and Pitts, 1943; Hebb, 1949; Rosenblatt, 1958] were a step in this direction, but at the time those ideas were formed, knowledge of the brain was much less than it is now, and the abstract nature of computation was less well understood. The major attraction of connectionism is that it can stand the crucial test of timing. That is, given that entire behavioral responses can be realized in a few hundred milliseconds, connectionist models of neural units seem to be the only way to achieve these response times. Previous papers have suggested how a particular connectionist theory of the brain can be used to produce testable, detailed models of interesting behaviors [Feldman and Ballard, 1982; Ballard, 1984; Feldman, 1981; Hinton et al., 1984]. The purpose of this paper is to relate these ideas more closely with emergent anatomical and physiological knowledge of the cerebral cortex.

2. Overview

Our principal hypothesis is that:

A major function of the cortex is to compute collections of *invariants* at different levels of abstraction.

An invariant is a description of a given situation in terms of a small number of parameters. For example, rigid motion can be described by six parameters: a rotation (three) about an instantaneous center (three). The usefulness of small-parameter descriptions that describe a large number of different situations pervades all aspects of perception and cognition. We suggest that the cortex has adopted representations and computational strategies that make the computation of invariants efficient.

2.1 Representing Possible Invariants

Any particular representation of information will have attendant advantages and disadvantages. Yet neurons in different parts of the brain seem to be representing information in very specific ways. One way to study these differences is through the single cell recording technique [Eccles, 1957] that allows recording of individual cell responses to different stimuli. Let us compare single unit recordings from neurons which function as basic parts of the ocular-motor system [Robinson, 1978] with those from cortical neurons in the visual areas [Hubel and Wiesel, 1962]. Within the ocular motor system, neurons with linear outputs are seen (Fig. 1a); that is, neurons whose *firing rate* is proportional to a scalar parameter, such as the rate of eye rotation. These neurons are part of the servo system controlling eye position and can be modeled as *summation* or *integration* devices. Their output has two important features. First, a larger value for their output variable means more frequent pulses. Second, the variable output is one-dimensional. In contrast, in the visual areas of cortex (and indeed all of the cortex), most neurons seem to be using fundamentally different encodings of their output (Fig. 1b, c). These neurons have *multi-dimensional receptive fields (RFs)*.¹ If the input stimulus is within its receptive field a neuron will fire; otherwise it is more or less quiescent. The degree of match between the stimulus and receptive field determines its firing rate. In the physiologist's terms, the first kind of neuron uses *frequency coding*, and the second kind of neuron uses *spatial (place) coding*. We have termed units using the latter kind of encoding *value units*. As first suggested by Barlow [1972], the value unit way of representing information seems to be a property of most cortical cells. The characteristics are the representation of a portion of a complex stimulus that is signaled if the cell is firing.

Figure 1.

Value units are a general way of representing different kinds of *multi-dimensional variables and functions* without requiring that each unit have a large bandwidth. The most optimistic estimate of discriminable signal range in a single neural output is about 1:1000, and this is insufficient to handle multi-dimensional variables. Value units overcome this limitation on neural output by breaking the ranges of a variable up into intervals and representing each interval with a separate unit. These intervals can be organized in many different ways. One straightforward way is to represent a




Figure 1: Different encodings as indicated by neural outputs. f = firing rate; v = stimulus value (a) Monotonic firing rate increases with scalar parameter indicative of a variable neuron. (b)/(c) One- and two-dimensional receptive fields indicative of value neurons. Variable neurons are typically rapid-firing (500 Hz) whereas cortical neurons are slow firing (50 Hz). Thus highly schematized plots in (b) and (c) are representative of the average of many trials; in other words, the post-stimulus histogram.

variable $\mathbf{v} = (v_1, \dots, v_k)$ isotropically by allocating a unit for each of N^k discrete values. These values are the center of intervals of width $\Delta\mathbf{v} = (\Delta v_1, \dots, \Delta v_k)$. The value k is the *dimensionality* of the variable. We will use the term parameter to denote a component of a variable.

Another way of thinking about the value/variable distinction is to compare the value unit encoding with variable representations in conventional, Von Neumann computers. In a Von Neumann machine, variables only access one value at any instant, and acquire these values by assignment statements. For example, $x := 3; y := 4$ assigns values 3 and 4 to x and y , respectively. Since a sequential computer can only access one value of a variable at a time, the notion of unique values for each variable at any instant is particularly appropriate. However, a parallel computer typically requires access to many values of a variable at the same time, and thus requires a different encoding scheme. A value unit representation such as an array of possible (x,y) values allows this parallel access. This difference is shown in Figure 2a and b. Two other examples of value encodings are also shown in Figure 2. Figure 2c shows a highly stylized representation of the orientation-sensitive cortical neurons found in striate cortex. In this case the variable is three-dimensional, with $\mathbf{v} = (x, y, \theta)$. Each unit represents a specific value of (x_0, y_0, θ_0) and has an associated $(\Delta x, \Delta y, \Delta\theta)$ that may be loosely thought of as its receptive field. Figure 2d shows a value unit representation for unit directions in three-dimensional space. In general, the intervals of neighboring value units will be overlapping.

Figure 2.


An important advantage of the value unit organization is that complex functions can be easily constructed. For example, suppose one neuron is sensitive to a red surround and green center and another is motion sensitive. In this hypothetical example, it is easy to see that by combining these two neuronal inputs at a third, one could construct a neuron that was sensitive to moving, red surround, green center input. Similarly one can combine responses to edge- and movement-sensitive neurons to construct more complex responses. This strategy can be used in a general way to construct arbitrary functions. Suppose one has such a function $f(x,y)$. Let us allot neurons for each interesting *value* of x and have a similar set for the interesting y values. One can think of these different values as very similar to just-noticeable-differences. Then these neurons can be used pairwise to construct the function as follows. The pairs of connections make synaptic connections with neurons representing appropriate values of f . We assume that both members of a pair of connections must be firing before the unit representing a specific value of f will fire. This type of input has been termed a *conjunctive connection* [Feldman and Ballard, 1982] and could be realized by an appropriate spatio-temporal summation behavior of multiple synapses. The main point, however, is that using this table look-up strategy, arbitrary functions, e.g., $f(x,y) = e^{\sqrt{x} \cdot \pi^{1/y}}$, are easily represented.

The ability to represent multi-dimensional variables and functions is especially important in vision. At early levels of visual processing, the useful information about complicated stimuli is only implicit and is generally distributed over space. For example, in navigation, the information needed to characterize one's trajectory stems


$X := 3$

$Y := 4$


a



b



c



d

Figure 2: (a) Von Neumann encoding assigns a single value to a variable at any instant. (b) In contrast, value unit encoding allows many values of the same variable to be accessed simultaneously. Two more examples illustrating the discreteness of value unit encoding are shown by (c) and (d). (c) Three-dimensional x-y-orientation units for striate neurons. In this scheme retinotopic coordinates and orientation are regarded as a three-dimensional space that is covered by the receptive fields of units. (Note that this space is actually toroidal, owing to the periodicity of θ .) (d) Two-dimensional units for indicating directions in space. Directions can be described by coordinates on the unit sphere. The sphere can be tessellated (covered) with uniform intervals. The figure shows a triangular tessellation based on subdividing an icosahedron. If one imagines the figure as a kind of helmet, then the units can indicate directions in space with respect to the head.

from change measurements distributed over disparate parts of the retina. The value representation for cortical cells allows the requisite trajectory parameters to be computed by transforming the input through a succession of levels of abstraction, each level being represented by value units [Ballard and Kimball, 1983]. This organization solves the problem without directly interconnecting outputs from all combinations of disparate sensors. This strategy requires N^2 connections between N sensors. Such a solution may be realized by insects, owing to their smaller numbers of retinal and cortical cells [Olberg 1981a; 1981b]. However, for animals with huge numbers of retinal and cortical cells, an easier solution is to introduce intermediate levels of value units.

The value unit concept is reminiscent of the idea of a "grandmother cell." Thus it could suffer from one of the principal criticisms of grandmother cells: if the cell dies, the function is lost. Networks of value units can be designed, however, that are impervious to such damage by: (1) distributing the function among small groups of local units; and (2) allowing for the recruitment [Feldman, 1981] of new, previously unallocated units. Since these notions are merely refinements to the useful idea of a value unit, we will keep referring to units as if they were represented by single neurons.

How can one determine if the value unit encoding is being used? The basic requirement was described in Figures 1b and c, but can be elaborated as shown in Figure 3. To determine value unit encoding, the peak responses of the units must span the useful range of measured values. The "useful range" depends on the role of the variable, but may be deduced by various means. For example, orientation value units should span the range $[0, 360^\circ]$. Thus, Figure 3a shows a variable that is *not* value encoded, whereas Figure 3b shows a value unit encoding of a variable. For our purposes, it is not necessary to have *all* the dimensions of a variable be value encoded. There are much stronger predictions made by the value unit encoding hypothesis, namely in the interconnection patterns that would appear between different cortical areas. These predictions require some development, and are discussed in Section 4.

Figure 3.

One must keep in mind that the value unit model is chosen as the simplest useful level of abstraction of the cortical neuron. Shaw [Shaw et al., 1982] has proposed that about thirty neurons are usefully considered as forming a functional unit. Abeles [1982] considers several interesting firing properties that arise from analyzing detailed recordings of interesting simultaneously-recorded units. One is that a neuron is far more likely to fire if its inputs are synchronized. Another very important idea is that neurons may be signaling probabilities [Hinton and Sejnowski, 1983]. The value unit model is not incompatible with these more detailed models, and is the simplest description that can illustrate important computational properties.

The notion of the neuron as a functional unit dates from the very earliest anatomical studies [Golgi, 1879; Cajal, 1911]. By way of contrast, other investigators have suggested that patches of dendrites might act as functional units [Shepherd et al., 1985]. It could turn out that the value unit model is extendible to dendritic trees





Figure 3: Firing rate, f , versus the value of a parameter, v , for neurons exhibiting different encodings. (a) Ensemble of neurons that are variable-encoding exhibit monotonic behaviors. (b) In contrast, value encoded neurons exhibit response maxima that span the measurement space.

considered as separate units. For purposes of explanation, however, we will stick to the level of abstraction that considers a neuron as a unit.

To summarize, neural responses can be used in two qualitatively different ways, which we have termed value encoding and variable encoding. The brain may have both value and variable neurons or use combined strategies, but the gross organization of the cortex seems to exhibit value encoding. Table 1 shows the main points of difference. Conversions between these two kinds of representations are well within known synaptic properties. These two representations can be combined in a hybrid system with the value neurons computing complex functions and the variable neurons computing dynamic control signals.

Table 1: Neurons with the hypothesized different encodings have fundamentally different characteristics.

<u>Value Encoding</u>	<u>Variable Encoding</u>
multi-dimensional representation	single-dimensional representation
-----	-----
receptive-field response with represented parameter	monotonic response with represented parameter
-----	-----
peak responses span parameter values	peak responses may not span parameter values

2.2 Computing Specific Invariants

To satisfy our criteria for a cortical model we must not only show how value units can represent information, but also how to compute using this representation. A strong argument for value units is the facility with which they support parallel computation. The reason for our emphasis on parallelism is that a growing amount of psychological data suggests that many complex behavioral responses require only a few hundred milliseconds (e.g., [Treisman and Gelade, 1980]). Parallel processing is indicated in these experiments since the response time is independent of the number of tokens that must be processed. In addition to these experiments, an elementary analysis of the neural signal also suggests that sequential processing is unlikely.

To illustrate how value units can solve problems in parallel, we will describe the solution to a very specific problem. Consider the simple map in Figure 4a, with four regions. The problem is to color the map so that no two adjacent countries have the same color. Each region may be colored with one of the colors shown. This problem is representative of a ubiquitous class of problems which can be posed as: "satisfy the largest set of compatible constraints" [Freuder, 1978; Hummel and Zucker, 1983; Prager, 1980; Rosenfeld et al., 1976; Ullman, 1979]. When this problem is translated to value unit notation, the color of each region is a separate value unit. If a particular

color is compatible with the currently chosen value units representing neighboring colors, then that unit is likely to be chosen to represent its corresponding region's color. The constraints are represented as links between units. There are many different ways to do this. We choose to let connections between locally incompatible colors be inhibitory (lower confidences) and connections between compatible colors be excitatory (raise confidences). These links are shown in Figure 4b. For brevity, two symmetric links are drawn as a single double-ended link.

Figure 4.

Now we shall describe how networks of value units compute. One can think of the i^{th} unit as having a small amount of information, (s_i, w_i) , where s_i is the state and $w_i = \{w_{i1} \dots w_{in}\}$ is the synaptic weight vector. In the underlying computational model [Hopfield, 1982] that we adopt, value units can be thought of as binary threshold units. Units start out in an initial state and converge to a stable final state. The state changes according to

$$s_i := 1 \text{ if } p_i \geq 0, \text{ else } s_i := 0$$

where

$$p_i = \sum_k w_{ik} s_k$$

This algorithm is only guaranteed to find local minima, but a feature of certain constraint satisfaction problems is that a local minimum is sufficient for the task at hand. For example, if the map coloring example is started in an arbitrary state, it will quickly converge to the state representing compatible colors, as shown in Figure 4. In this case, the negative weights = -2, and the positive weights = 1.

An extension of this algorithm has been made by [Hinton and Sejnowski, 1983], building on the concept of "simulated annealing" developed by [Kirkpatrick et al., 1983]. In this algorithm, the state is adjusted probabilistically according to

$$s_i := 1 \text{ with probability } P_i$$

where

$$P_i = 1 / (1 + e^{-p_i/T})$$

In this updating rule the parameter T plays the role of "temperature" (analogous to its role in a Boltzman distribution). The temperature is initially high, corresponding to equally probable state changes, and is gradually "cooled." The advantage of this algorithm is that it now finds a *global minimum* with a given probability. A drawback of the probabilistic version is that it may require too many iterations to be biologically plausible, but the discovery of this algorithm is very recent, and ways may be found to overcome this problem. Also, although the algorithm formally requires symmetric weights, computer simulations have shown that this condition




Figure 4: (a) A four-country map coloring problem: color the map with the colors shown (r = red, g = green, y = yellow) so that adjacent countries have different colors. (b) A value unit representation of the map coloring problem. A separate unit is allocated for each color of every country. Symmetric inhibitory connections are denoted by double-ended links terminating in small circles. Symmetric excitatory links are shown by double-ended links terminating in arrows. Other designs are possible. For example, one could use inhibitory connections between incompatible colors in neighboring countries. The particular design keeps inhibition local to individual countries. (c) A particular starting configuration where dark units are on and light units are off. (d) The correct solution achieved by iteration using the state modification rule described in the text.

might be able to be relaxed. Important details of this model are given in [Kirkpatrick et al., 1983; Geman and Geman, 1983; 1985; and Hinton et al., 1984]. In particular, the second of these papers addresses convergence properties, and the last provides a learning algorithm. Our own extension, elaborated in later examples, uses non-symmetric *ternary* weights (w_{ijk}) that model spatio-temporal summation (which we have called conjunctive connections).

The points behind the map coloring problem are threefold. First, the problem is characteristic of other constraint satisfaction problems in that empirical tests show that larger-scale versions do not require more time. Maps with more countries and colors can still be colored without any increase in processing time if the constraints are appropriately organized. This controversial statement about problem scaling is currently based on empirical tests. Kirkpatrick et al. argue that convergence is based on how "frustrating" (incompatible) the constants are. We will take up this point again in the context of the specific examples in Section 5. The second point is that the kinds of constraints that we used are extremely general and can characterize a broad range of perceptual and cognitive situations [Hinton et al., 1984; Feldman, 1982; Ballard and Hayes, 1984]. In particular, problems in visual gestalt recognition can be described as trying to satisfy an appropriately weighted collection of local constraints [Ballard et al., 1983; Feldman, 1982]. The third lesson of constraint satisfaction is that *local* constraints can imply a *global* solution.

There is a sense in which this problem can be thought of as an analogy to how more complicated problems are handled. If one interprets "map" as "cortical area," one sees how different local constraints might participate in a global computation. The requirement for a unique color for each region would be analogous to an inter-area cortical constraint, whereas the constraint between neighboring colors would correspond to an intra-area constraint between different cortical areas. However, to elaborate on this point: it would obviously be a big mistake to think that this kind of computation is *all* that the cortex does. Our view is that a large class of problems crossing many domains (e.g., vision, motor control, portions of cognition) can be potentially solved in this fashion.

It is far from established that the cortex actually uses its signals in a way suggested by the constraint satisfaction paradigm, although, as Terry Sejnowski suggests, the use of the post-stimulus histogram in data presentation implicitly appeals to an underlying statistical model. To go further, one might try to: (a) establish the role of a given neuron in a specific visual task; and (b) see if its firing rate increases or decreases markedly during the few-hundred-millisecond convergence period. With current techniques, this would be a very difficult experiment to do. Several experiments show that the response of cortical neurons is dependent on the task that the animal is engaged in (e.g., [Shaw et al., 1983]), but so far these are still not definitive in implicating the constraint satisfaction paradigm. On the other hand, there is no evidence to *rule out* the constraint satisfaction model, either. One problem is that very few alternative testable models have been proposed that meet the computational requirements. Holographic memory has been proposed, but this has the form of static memory and does not address the problem-solving issue.

1Traditionally, the receptive field of a neuron has been defined with respect to a stimulus. This has the unfortunate effect of making it dependent on a particular experiment. In our model we define the receptive field as the response to *all* the neurons' inputs, some of which may be feedback connections from neurons in more abstract cortical areas.

3. A Review of Cortical Organization

The value unit model has several implications for cortical organization, but in order to develop these implications, we first need to summarize several relevant cortical anatomical and physiological features. The summary deserves the following cautionary note. It represents an attempt by the author to highlight important neurobiological features that are relevant to the proposed model. As such, it may deter neuroscientists from appreciating the abstract computational components of the model which are aided but not vitally tied to the specific neurobiological associations. The other side of this is that committing the model to a particular anatomical substrate makes it easier to understand its salient features.

3.1 Cortical Columns: A Uniform Processing Architecture

To a coarse approximation, the cortex can be regarded as a two-dimensional sheet, a few millimeters in thickness. Within this sheet the anatomy can be usefully thought of as being organized into functional, overlapping columns of about 800 μm in diameter [Mountcastle, 1978 (page 21)]. Metaphorically, one can think of tightly packed interpenetrating cylinders stacked axis-parallel, where the diameter of each is .3 mm and the length is a few millimeters. Szentagothai [1978a; 1978b] has done extensive work that also shows elaborate anatomical similarities in structure of cortical columns from area to area and many others have confirmed the columnar organization into interdigitated, repeated functional units [Goldman-Rakic and Schwartz, 1982; Kaas et al., 1981; Lund, 1981].

Superimposed on this organization is another organization dictated by the interconnections between the cortex and other parts of the brain. For this purpose, the sheet-like cortex may be thought of as approximately divided into three major layers. The uppermost layer (sublayers II and III) contains neurons that are connected to other neurons in different parts of the cortex. The middle layer (layer IV) receives input from other parts of the brain, and the bottom layer (V and VI) contains neurons which handle cortical output to other parts of the brain.

Our main interest is in the connection patterns between different neuronal units in functionally different cortical areas. Recent experiments strongly suggest that the axonal arborization is an important vehicle by which the neuron expresses functional diversity. For example, three-dimensional reconstructions of striate neurons stained with HRP show that the axons exhibit striking geometric orientation preferences [Gilbert and Wiesel, 1983]. As another example, Blasdel [Blasdel et al., 1983] has shown a case in striate cortex where the axonal buttons from a neuron in layer IV remain exactly within an ocular dominance column. This is only one experiment, but its important implication is that the connections are an important way to achieve functional diversity.

3.2 Functionally Different Cortical Areas

The two-dimensional, six-layer structure holds for almost all of the cortex, but at a lower level of detail the cortex can be differentiated into distinct areas. In fact, it has been known for a long time that the cortex itself can be divided into different cytoarchitectural regions [Brodmann, 1909], with the different regions almost always

having different functional characteristics. These functional characteristics have been determined through a variety of techniques, including anatomical track-tracing, lesion studies, radioisotope labeling, and single electrode recordings (e.g., [Zeki, 1978]). Figure 5 shows different areas for owl monkey cortex after Allman [Allman et al., 1982].

Figure 5.

The figure shows predominantly visual areas and somatosensory areas. Other areas are in the process of being mapped. The visual areas are all *retinotopic*; that is, moving an electrode across the cortex in a visual area corresponds to traversing the visual field in a retinal coordinate frame. The different visual areas each exhibit geometric distortions of the visual field as hinted at by the different midline locations shown on the figure. A helpful method of identifying such fields is to stain callosal fibers connecting the two hemispheres, since companion visual areas are densely interconnected at the regions corresponding to the midline in the visual field [Van Essen et al., 1982].

The first visual area, V1, seems to represent primary visual parameters such as disparity, orientation, luminance, change, and color. Other areas seem to represent more abstract features. For example, MT neurons have been found that are sensitive to changes in physical motion.

The visual areas can be coarsely characterized as to whether they seem to be involved in stabilized vision or motion. For example, with reference to the owl monkey, DL and DM [Allman et al., 1982], which have expanded foveal representations of the visual field, would seem to be carrying out computations important in stabilized vision. In contrast, MT and M, which are most responsive to moving random dot patterns, would seem to be carrying out computations important for motion.

Areas adjacent to the visual areas generally represent very abstract parameters. However, characteristics of the visually responsive neurons in superior temporal sulcus are that they are responsive to non-retinotopic stimuli [Gross et al., 1982]. For example, Sakata [Sakata et al., 1980] has identified neurons responsive to full-field rotations in parietal cortex.

3.3 Abstraction Hierarchies

Different cortical areas seem to *represent information at different levels of abstraction* [Van Essen and Maunsell, 1983]. As an example we consider the relation between intensity changes and optic flow. Optic flow is a retinotopic projection of the three-dimensional velocity field. Early in the visual areas neuron RFs are sensitive to some kind of motion, but there is a distinction between motion as reflected in time-varying intensities and motion as represented by optic flow: optic flow neurons respond to physical velocity changes and not to illumination changes.

An important experiment is that of [Movshon, 1983], which compared the response of neurons in V1 and MT in the macaque. Given a checkerboard stimulus,

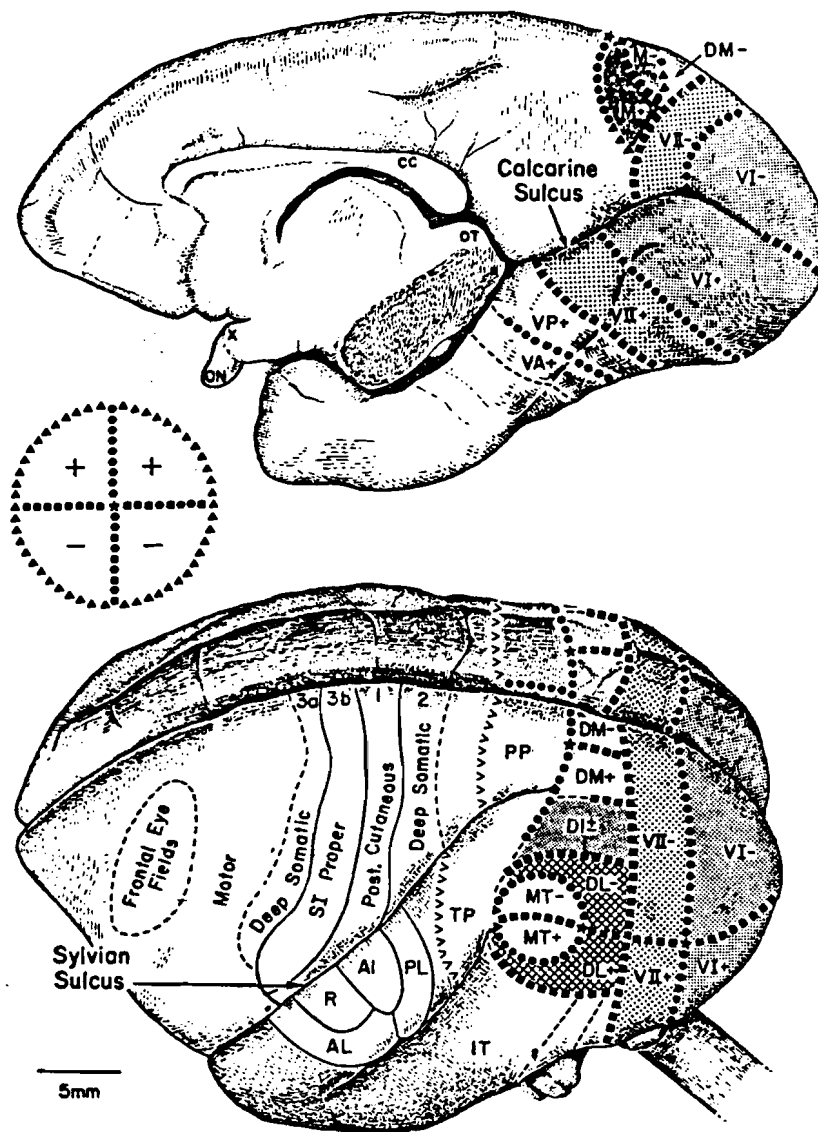


Figure 5: The representation of the distinct sensory domains in the cerebral cortex of the owl monkey. Above is a ventromedial view of the right hemisphere; below is a dorsolateral view. DI, Dorsointermediate Visual Area; DL, Dorsolateral Crescent Visual Area; DM, Dorsomedial Visual Area; IT, inferotemporal Cortex; M, Medial Visual Area; MT, Middle Temporal Visual Area; PP, Posterior Parietal Cortex; VA, Ventral Anterior Visual Area; VP, Ventral Posterior Visual Area; AI, First Auditory Area; AL, Anterolateral Auditory Area; CC, Corpus Callosum; ON, Optic Nerve, OT, Optic Tectum; PL, Posterolateral Auditory Areas; R, Rostral Auditory Area. From [Allman et al., 1982].

neurons in V1 responded optimally when the motion was perpendicular to the intensity gradients of the checkers. Some neurons in MT responded when the motion of the checkerboard was aligned with their preferred direction. While the number of neurons tested in MT was small, the result hints that the representations that are arrived at by abstract mathematical analysis of the physical constraints may be directly implemented in the underlying neural structure. Furthermore, the neural responses are compatible with value unit criteria (Table 1).

Anatomical evidence that cortical areas may represent information at different levels of abstraction comes from the consideration of cortico-cortical connection patterns [Maunsell and Van Essen, 1982]. If area A is more abstract than area B, it will receive different connections to its layer IV and send connections to B's layers II, III, and V, or II, III, and VI. Using this result and the connection patterns for the owl monkey [Allman et al., 1982], one can construct the following hierarchy (Fig. 6). This hierarchy also can be partitioned into form (DL-DM) and motion (MT-M) channels, following Maunsell and Van Essen's observations on the macaque. The difference between form and motion is that form requires the processing of distributed, spatial locations, whereas motion (in the case of a rigid body) can be summarized as a few *feature* values such as rotation and translation. Evidence for the form-feature distinction comes from a number of experiments (e.g., [Mishkin et al., 1983]). These show that macaques with inferotemporal lesions perform poorly at feature recognition tasks whereas macaques with parietal lesions perform poorly at spatial location tasks.

Figure 6.

Much data exists to show that orderly abstraction hierarchies are also present within an individual area. For example, Hubel and Wiesel [1962; Hubel et al., 1978] described simple, complex, and hypercomplex neurons in striate cortex, and other functional experiments imply that the most abstract neurons in an area reside in the upper layers [Movshon, 1983; Pasternak et al., 1981].

Although the definitive experiments have not been done, preliminary evidence from [Woolsey people] suggests that similar hierarchical organizations exist in motor and somatosensory areas.

3.4 Spatio-Temporal Channels

Another important feature of cortical structure is the retino-cortical pathways. The neurons in these pathways have markedly different spatio-temporal responses (for a review, see [Stone et al., 1979]). Neurons in these pathways can be broadly classified into three types:

- 1) X-cells, which have fine spatial resolution, but coarse temporal resolution;
- 2) Y-cells, which have coarse spatial resolution and fine temporal resolution; and
- 3) W-cells, which have relatively coarse resolution in both space and time.

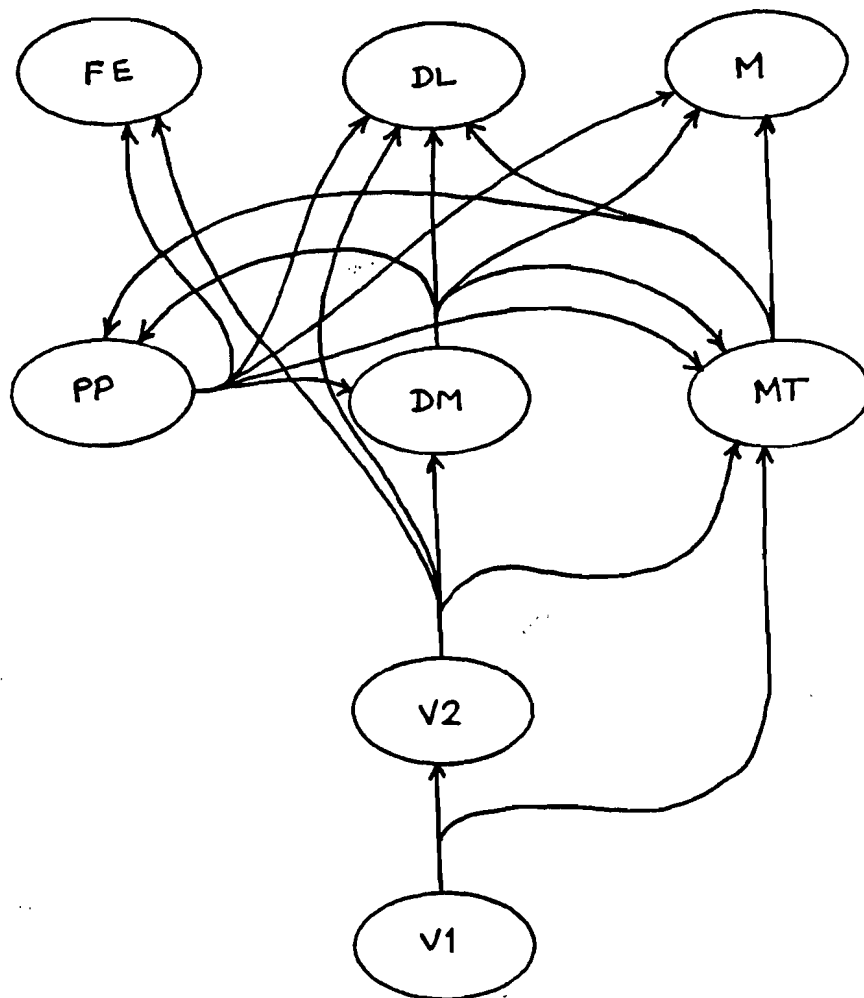


Figure 6: Organization of owl monkey cortical areas into a functional hierarchy using Maunsell and Van Essen's [1982] connection principles and data from Weller and Kaas [1982]. Arrows denote direction of hierarchy using either the feed forward or feedback principle.

The X-Y-W system is involved in two distinct spatio-temporal channels that are distinguished by involving different subcortical nuclei and by different cortical afferents. The most studied pathway is from the retina through the lateral geniculate to striate cortex. In the cat, the X-Y-W systems have separate locales in both the LGN and cortex, and this result seems to be true also for monkeys.

The other pathway is from the retina to the superior colliculus, to the inferior pulvinar, and then to cortex. The interesting results here are: (1) only the Y-W neurons seem to be involved in this pathway; and (2) the cortical afferents terminate in most retinotopic areas [Weller and Kaas, 1982]. These two channels are summarized in Figure 7.

Figure 7.

The discussion in this section has been focused on the cortex in order to describe findings that relate to connectionist models. The main points are:

- 1) The processing architecture is surprisingly uniform when the cortex is considered as a two-dimensional sheet of layered processing columns.
- 2) The two-dimensional sheet is divided into different functional areas; that is, areas wherein single unit recordings reveal different and characteristic responses.
- 3) Emergent evidence suggests that much of the functional diversity is realized through axonal connection patterns.
- 4) The different areas form a natural hierarchy which is revealed through the layered patterns of efferent and afferent connections.
- 5) The hierarchies can be thought of as further divided into channels. The principal channels are spatio-temporal channels and have direct anatomical correlates in terms of cell types.




Figure 7. Spatio-temporal channels. The two principal pathways to cortex: (1) the tectopulvinar relay system (Y-W only) and (2) the lateral geniculate relay system (X-Y-W). After Weller and Kaas [1982].

4. Value Unit Constraints on Cortical Organization

One of the most interesting aspects of the value unit model is that it has huge consequences for cortical organization. This section shows how the value unit representation constrains the geometry of the cortical layout.

4.1 Relating Function to Cortical Topography in Striate Cortex

Let us first consider striate cortex. One of its main characteristics is an explicit representation of the visual environment. Recently, Marr [1978] and Barlow [1981] have articulated the notion of explicitness in cortical representations. An old puzzle was the huge ratio of cortical to retinal neurons, since in some sense the cortex cannot add to the information represented in the retina. Most data suggests that while the cortex does not expand retinal information, it does make *implicit* information *explicit*. The visual field is mapped onto it retinotopically. The entire visual field is represented, but the central portion is magnified by having more cortical cells. Neurons in a certain area will only respond to inputs in a particular part of space, but within that cortical area several other parameters are represented. There are neurons sensitive to edge orientation [Hubel and Wiesel, 1962], scale (or "spatial frequency") [De Valois, 1977], and ocular dominance [Hubel and Wiesel, 1962], as well as other parameters for a particular part of visual space. Tootell [Tootell et al., 1983] has shown that the representations of these different parameters occupy different but overlapping areas of cortical space. Since *all* values of this list of parameters may occur at a particular point in space (and its corresponding locale in the cortex), they must be represented somehow within that small cortical space. One way of doing this would be to have a neuron for discrete values of every parameter (e.g., each value for scale (spatial frequency), each value for orientation, etc.), and to a first approximation this seems to be the underlying representation. The major refinements are: (1) that neurons generally have multi-parameter responses, e.g., a neuron that is sensitive to orientation will generally be sensitive to certain values of velocity, as well as other parameters; and (2) neurons will have overlapping receptive fields.

4.2 Primary and Secondary Indices

The place-coded structure constrains the functional layout of striate cortex in ways that can be generalized. We argue that the problem is to represent multi-dimensional fundamental parameters in a two-dimensional architecture. This constraint usually means that two of the parameters can be regarded as the *primary indices*. Variations in these two parameters span the cortical area. In all the visual areas, the primary indices are retinal coordinates, and *secondary indices* are important visual parameters, such as orientation, colors, and motion. For each value of the primary indices, a complete set of secondary indices will be represented. Figure 8a shows a schematic of some of the parameters of striate cortex in the macaque using secondary indices of edge orientation and ocular dominance. This basic organization has been found for many other cortical areas. For example, the body is represented four different times in Brodman's areas 1, 2, 3a, and 3b [Kaas et al., 1981]. The primary indices are topographically related to the body surfaces and the secondary indices are related to tactile and joint sensory parameters. Sensory areas have been found where the primary indices are not space (e.g., auditory cortex, olfactory cortex)

but, in general, the correct non-spatial primary indices are difficult to deduce, since one must guess the correct parameterization. An area where the correct guesses may have been made is the frontal eye fields [Goldberg, 1982] which govern eye movement. Here the primary indices seem to be: the magnitude and direction of the next eye movement represented as a vector. Secondary indices that have been discovered are the last eye movement, also represented in terms of a radius, direction. Figure 8b schematizes this organization, and Table 2 summarizes some current physiological findings in terms of the primary and secondary indices, drawing on data from [Maunsell and Van Essen, 1982; Tusa et al., 1979; Tusa and Palmer, 1980; Burton and Robinson, 1981; Juliano et al., 1983].

Figure 8.

Although the primary indices for the somatosensory areas are listed as being body-topographic, this classification is becoming more tentative. Recent evidence [Burton and Robinson, 1981; Juliano et al., 1983] shows that the body parts are multiply represented. This may be due to the fact that body topography is a secondary index in these areas, and that there exists a primary index with a broader classification. One suggestion is that a body part is organized by broad kinds of *tasks* in which it participates. Thus when a finger is used alone it might be in one locale, whereas when it is used with digits from the opposite hand, it may be in another locale (but both would be regarded as being within a task-indexed cortical area).

Many neurons, for example frontal eye field neurons, have large overlapping RFs, and this brings up an important point. The indices are defined by the parameter value that gives the maximum response. The neurons may have large RF widths that are related to their *inputs* and are, of course, totally unrelated to the width of their cortical columns.


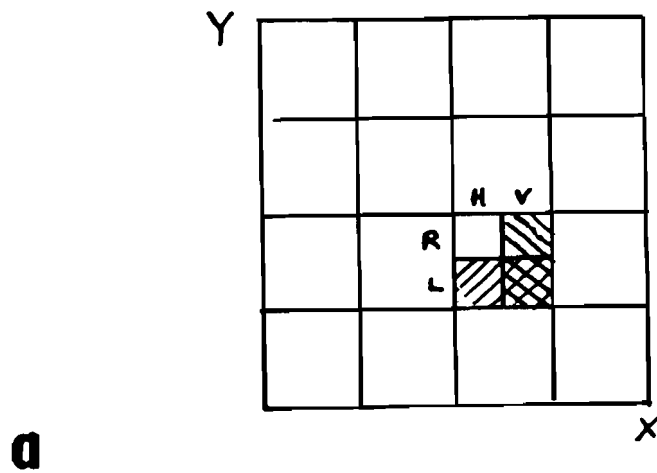


Figure 8: (a) Schematic of organization of striate cortex. Primary indices are X and Y; two of the secondary indices are ocular dominance and orientation. (b) Schematic organization of frontal eye fields (after [Goldberg, 1982]). Primary indices are next eye movement R, ϕ . Secondary indices are last eye movement r, θ .

Table 2.

MONKEY (macaque)

Area	Primary Indices	Secondary Indices
V1	space (retinal)	spatial frequency, color, orientation, ocular dominance
V2	space (retinal)	motion, disparity
V3	space (retinal)	
V3A	space (retinal)	
V4	space (retinal)	color
MT	space (retinal)	motion
FEF	Δ space (next eye movemt.)	Δ space (last eye movemt.)
1, 3b	body surface	cutaneous parameters
3a, 2	topological space	joint & muscle parameters

CAT

Area	Primary	Secondary
17	space (retinal)	spatial frequency, ocular dominance, orientation
18	space (retinal)	motion
19	space (retinal)	

The primary/secondary index dichotomy is meant to be a useful distinction to aid in the interpretation of the possible function of cortical areas. An important refinement that this distinction does *not* address is the details of the packing of secondary indices. Two recent models for striate cortex [Dow and Bauer, 1983; Cynader et al., 1983] directly address this issue, as does work by Tootell [Tootell et al., 1983], Livingstone and Hubel [1984], Hubel and Wiesel [1963], and Singer [1981].

4.3 Space Limitations Require Different Functional Areas

We term a regular organization of neurons in the cortex, like those in V1, a *parameter net*. Its characteristics are that all values of a small set of parameters must be represented within a certain area. One important question is: why should these parameters be clustered in this fashion [Cowey, 1981]? Our answer is that, in any given area, a severe packing constraint follows from the fact that neurons have *multi-modal responses*. Such responses are exemplified by orientation-sensitive neurons which will also respond to motion, spatial frequency, and other stimulus variations. A

consequence of this property is that only about 5-15 scalar parameters can be represented in a given cortical area. This is because the number of units required grows exponentially with the number of parameters. If N is the number of just noticeable differences in each scalar parameter and k is the dimensionality of the unit (number of parameters), the total number of units required is N^k . Thus we refer to this severe packing constraint as *the N^k problem*.

Overlapping the receptive fields can alleviate the N^k problem, but not to a significant extent. The main benefit of overlapping, multi-dimensional units is that in an architecture where units are expensive and connections are relatively cheap, these units can allow for the representation of a given signal resolution with less units. Basically, one can *intersect* the RF of overlapping units through spatio-temporal summation, as shown in Figure 9. Hinton [1981] has shown that the savings is a factor of $1/D^{k-1}$ where D is the diameter of the receptive field, in units of desired maximum resolution, and k is the dimension of the stimulus.

The overlapping RF encoding scheme has an accompanying disadvantage. It cannot signal closely spaced stimuli simultaneously without error, as shown in Figure 9. Thus the price paid for the encoding economy is some loss in parallelism.

Figure 9.

It is informative to plot the number of units required as a function of N and k , as shown in Figure 10. The different curves correspond to different values of D . These curves support the assertion that the dimensionality k that can be represented in any cortical region must be on the order of 5-15, even given the overlapping field strategy. Thus a consequence of overlapping, explicit representations and the number of neurons available is a fundamental limitation on the number of parameters that can be represented in a given cortical area. If this is the case, then a natural consequence of having more parameters is to have more cortical areas. This may be one of the reasons for the several visual areas seen in the cat and primate cortex. Extensive data from electrophysiological recordings reveals that neurons in the different visual areas have different responses. Figure 11 shows data from Allman et al. [1982] for different cortical regions in the owl monkey. While all the visual areas have some neurons that respond to almost any given stimulus, certain areas have large portions of their neurons dedicated to particular subsets of visual parameters. Intriguingly, these parameters are similar to parameters that have been shown to be computable by parallel algorithms: color, edge orientation, disparity, surface orientation, and optic flow [Brady, 1982; Marr, 1982; Ballard et al., 1983].

Since low-dimensional spaces are an economical representation, one might wonder why minimal dimensions are not used. The answer is that low-dimensional spaces have special representational problems termed "illusory conjunctions." These are described in Sections 5 and 6.

Figure 10.

Figure 11.




Figure 9: (a) Non-overlapping encoding for a parameter space with $N=3$ and $k=2$. In general, the total units required by this scheme is proportional to N^k . (b) A coarse coding of this space uses three arrays of units (one unit from each array shown), where each unit has a diameter $D=3$. The diameter is defined in terms of the high resolution units. The simultaneous firing of all three units denotes input in area indicated. In general, the total units required by this scheme is N^k/D^{k-1} . (c) Developing the example in (b) further to show how closely spaced stimuli can lead to errors: stimuli at the shaded positions produce two errors denoted by Xs.




Figure 10: The fundamental restriction on the dimensionality of a cortical area. The graph shows the total number of units T required as a function of different values of k, N and D where $T = N^k/D^{k-1}$. Since the number of units in an area must be less than 10^{10} , the dimensionality of an area must be less than 15.

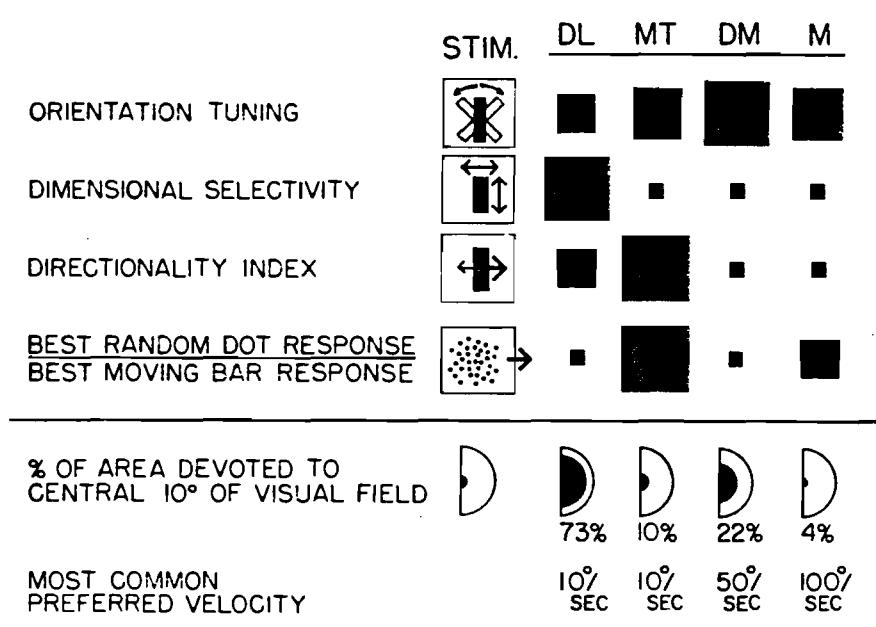


Figure 11: Functional specificity in visual areas DL, MT, DM, and M in the owl monkey. (Refer to Fig. 5 for the locus of these areas.) The strength of the functional attribute is indicated by the size of the black squares. From [Allman et al., 1982].

4.4 Effects of Changing Indexing Schemes on Anatomical Connections

If different cortical areas use value units but with different parameters, it is natural to ask how they are interconnected. A major consequence of the value unit hypothesis is interconnections between the different areas should show up as distinctive patterns that can be revealed by current histochemical staining techniques. Basically, retrograde staining between areas with different primary indices should produce distinctive spot-like or band-like patterns. Furthermore, these patterns should be related to functional organization. Evidence for such bands has been demonstrated in retrograde horse-radish peroxidase (HRP) staining [Gilbert, 1982; Montero, 1981; Curcio and Harting, 1978]. To repeat the important point, such patterns would arise from interconnections *between areas with different primary indices*. The different visual areas do *not* meet this condition, being indexed by space in every case. In fact, experiments by Montero [1981] on retinotopic areas in the cat strongly suggest that they are connected on a point-to-point basis.

To develop the argument for connections between areas where the primary index changes, we first consider an abstract example. Figure 12 shows a case of two representations which differ only in that the order of indices have been reversed. In Figure 12a the primary indices are (x,y) ; in Figure 12b the primary indices are (α,β) . Now consider the connections from area B to area A. In particular, consider all connections that synapse in the shaded box in A. This box contains units which have all values of α and β and a particular value of x and y . Where are these corresponding units represented in area B? One quickly sees that the global pattern of area A must be repeated for every column of different values of (α,β) in B. Thus retrograde staining from A to B would reveal these neurons as spots of stained cell bodies. Since the connections between the two areas are symmetrical, retrograde staining from B to A would reveal the same kind of pattern. An interesting feature of this organization is the insensitivity of the demonstration to the size of the site. Suppose, rather than a single cortical column, one considered several adjacent columns. Analysis shows that the spots in B are now larger but that the overall pattern is unchanged.

Figure 12.

Even when the indices change and are not α,β but instead $f(\alpha,\beta),g(\alpha,\beta)$, the same kind of patterns can occur. This analysis is derived from the interconnections between two different cortical areas. To expand on this point, consider the example of computing a global visual parameter with value units, say egomotion parameters, from a retinotopic map such as optic flow. Figure 12c and d show representative connections for the subset of value units that have zero rotation. Thus, wherever the primary indices in the two areas are different, the connections between areas should form a characteristic pattern indicative of a one-to-many mapping. Thus this model can explain the patchy areas of label in retrograde staining experiments [Gilbert, 1982; Goldman-Rakic and Schwartz, 1982; Curcio and Harting, 1978; Montero, 1981].

To understand our *intercortical* connection hypothesis it may help to compare it with the *intracortical* explanation advanced by Mitchison and Crick [1982] to explain


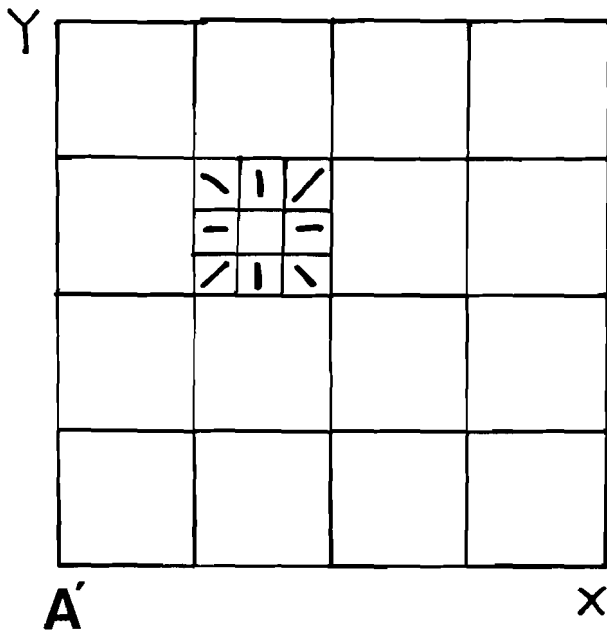
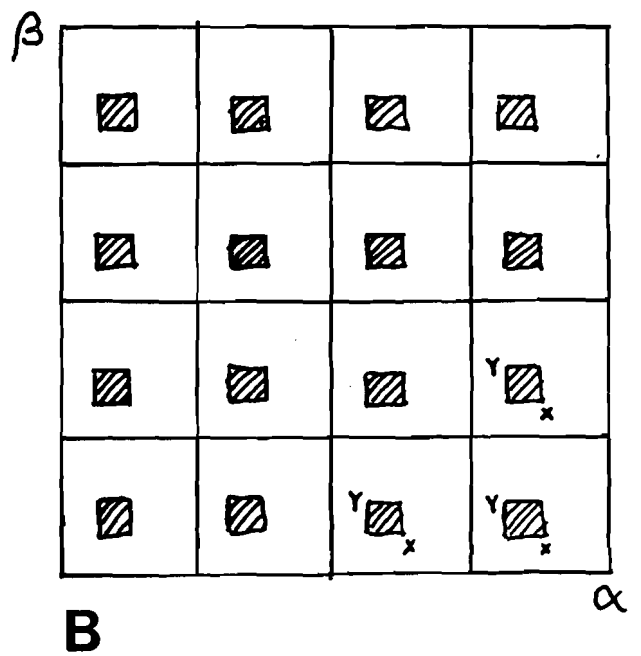



Figure 12: Two examples showing how connection patterns can change between retinotopic (A, A') and non-retinotopic areas (B, B'). The upper two figures show how connection patterns can change in a hypothetical situation where the primary and secondary indices are interchanged in two different areas. In Area A the primary indices are (X,Y) and the secondary indices are (α,β), and in Area B the primary indices are (α,β) and the secondary indices are (X,Y). The shaded areas are examples of those that would be connected on a point-to-point basis. The second example shows a subset of the connections in a situation where the retinotopic area A' represents optic flow as indicated by direction vectors and the non-retinotopic area B' represents rigid motion parameters for translation (U, V, W) and rotation. The primary indices are assumed to be U/W, V/W. Next we show a subset of the connections from A' to B' that represent values of rotation equal to zero. The three rotation parameters are not specified in the figure, but it is assumed that a rotation value of zero is at the center of each of the areas with different primary indices. Both of these examples are meant to be representative of the kinds of transformations that occur when the primary indices of different areas are different. The patterns may not be as regular as depicted here, but in general will be "point-to-many."

Rockland and Lund's [1982] tree shrew data. In the tree shrew, extra-cellular retrograde HRP transport from a striate cortex injection site produced patchy labeling within the striate area. Mitchison and Crick hypothesized that the patchy labeling arises from the orientation of long axons that connect neurons of similar orientation selectivity. In our terminology, the primary indices do not change in this case, but the particular secondary index that is labeled happens to depend on parameter values in a specific geometric direction. In short, the tree shrew data can be explained by connections between related parameters within the cortical area, but the same kinds of patterns could be produced as a consequence of changing indexing schemes in different cortical areas.

4.5 Developmental Constraints on Cortical Structure

The synopsis of the description of cortical representation of parameters was: a) that the cortical areas were arranged in patches, using primary and secondary indices for groups of 5-15 parameters; and b) interconnections between patches that had different primary indices should reveal characteristic spots or bands, whereas interconnections between patches with the same primary indices would be one-to-one. That this organization should hold for the entire cortex has been argued by extrapolating current experimental evidence. These arguments become weaker as the representation becomes more abstract. That is, numerical or topological parameter organizations are naturally represented with certain primary indices in the cortex. In many computations one can imagine, connections are needed to nearby values (speaking numerically). Matters are greatly simplified in a representation where the nearby numerical quantity is also an anatomical neighbor. This advantage lessens as the parameters become more abstract. In dealing with sets, one might expect to be able to shuffle the order of neurons (while preserving connections) without result. However, these are not the only constraints on cortical organization.

Another source of constraints is developmental: the brain has to connect itself up correctly in the first place. Experiments by [Rakic, 1974; 1981; Goldman-Rakic and Schwartz, 1982] have revealed an elegant pattern of cortical growth that produces *interleaved* connections between different areas. Interleaved connections may be achieved by the *temporal phasing* of growth [Rakic, 1981]. For example, connections from area A to cortical area B may initially be dense, but are spread apart by cortical growth. Another set of connections to B starting at a later time can use the first set of connections as guides to fit in between them in a uniform fashion. Thus the parameters of secondary indices may be naturally placed at the ends of appropriately interleaved connections to areas which use them.

To conclude Section 4, we summarize the main points:

- 1) The cortical architecture can be thought of as composed of functional areas, each representing a multi-dimensional parameter space.
- 2) Our hypothesis is that these parameters project onto the cortex in such a way that two dimensions are isomorphic to the cortical sheet and the rest are interleaved. The former are termed primary indices and the latter are termed secondary indices.

- 3) Constraints on unit value architecture suggest that only 5 to 15 parameters can be completely represented in any given cortical area.
- 4) Savings in units can be achieved with overlapping RFs but only with the sacrifice of parallel computing capability.
- 5) Connections between cortical areas with different primary indices are primarily one-to-many.
- 6) Value unit concept is compatible with current ideas about cortical growth whereby connections from different regions are temporally sequenced.

5. Computing with Cortico-Cortical Connections

The introductory section describes how value units can be used for a general class of computation. The map coloring example faithfully represents the abstract nature of the computation but is removed from general perceptual and cognitive problems. This section closes this gap by discussing two examples, from shape and motion perception, respectively.

The fundamental issue for value unit encoding is the N^k problem described in Section 4.3. Given this limitation, how can one represent information in a way that captures sufficient diversity? The answer is that there are many different strategies for decomposing high-dimensional structures into networks of value units, each of low dimensionality. However, two of these are the most important: hierarchies and subspaces.

Hierarchies split higher-dimensional spaces into *more abstract* and *less abstract* parts. As an example, consider some enumeration of tasks T , e.g., hammering, sawing, etc., with objects O that the tasks are applied to, e.g., nails, wood, etc. Here our notion is that a task is a more abstract concept than the object of the task. One could have units for each combination, i.e., $T \times O$, but a more economical decomposition is to have separate networks for T and O with constraints between them.

Subspaces are a mathematical concept that logically includes the above example, but we (somewhat arbitrarily) use the term for decompositions at the same level of abstraction. The idea is simple: if a value unit representation is infeasible for some dimension k , i.e., N^k exceeds 10^{10} , then split the space into m subspaces, each of lesser dimensions $k_1 \dots k_m$, such that

$$\sum_{i=1}^m N^{k_i} \ll 10^{10}$$

The first question that occurs is that of when the subspaces uniquely represent the original data. (This is very similar to the reconstruction from projections problem in computed tomography, except that now the units are either on or off.) Unfortunately, the answer depends on the number of high-dimensional units that would be *on*. If this number is n , then the number of subspaces of dimension k' ($= k_1 = k_2 = k_m$) that are required to unambiguously represent the high-dimensional data has been shown to be bounded by [Kemperman, 1982]:

$$1 + n((k \text{ div } m) - 1)$$

where div denotes integer division. This result requires the number of subspaces to grow as the complexity of the information. What this means is that, in general, some trade-off will have to be made. The subspaces can be chosen to unambiguously represent a certain number of concepts; after that point some ambiguity will be introduced into the representation. This ambiguity is dealt with in the next section.

The remaining parts of this section illustrate specific solutions to the N^k problem that use hierarchies and subspaces.

It is very important to understand the claim made by the following examples. Since these examples illustrate technical details of the value unit approach, they may also inadvertently overemphasize particular constraints compared to competing parameterizations cast in terms of conventional computing models. Better constraints may be found; the main point that these examples illustrate is that in order to be cast in terms of the value unit model, the constraints have to be expressed in terms of sets of relations, each involving only a small number of terms.


5.1 Shape Perception

One way of recognizing a familiar rigid shape relates visual features in a retinal coordinate frame to those of a stored prototype. The stored prototype is economically described with respect to an intrinsic frame which is related to some special features of the prototype itself (Figure 13a, b). To relate the retinotopic and intrinsic data, a coordinate transformation between the two frames must be computed. By choosing the visual features appropriately, this coordinate transformation can be easily computed if the correspondence between a prototype feature and retinotopic feature is known. We have shown that the correct transformation can be computed even when this correspondence is not known in advance: all possible correspondences compute transformation values, and the correct value is selected as the consensus among the different values computed.


One extremely plausible use of cortico-cortical connections is for computing these kinds of coordinate transformations. The key to this computation is to represent possible transformation parameters as explicit value units [Ballard, 1981; Hinton, 1981]. Once this is done, any transformation $x' = Tx$, where x' , x are data in two different coordinate frames and T is the transformation between them, can be captured as a relation between an explicit x' unit and explicit transform unit and an explicit x unit [Hinton, 1981]. Figure 13c shows this relationship diagrammatically.

Figure 13.


The diagram shows conjunctive connections (described in Section 2.1) to a representative transform unit. Thus, a T unit will only receive input if both the x' unit and the x unit are on. To make this more specific, consider the domain of two-dimensional polygonal shapes such as that shown in Figure 13a and b. Let each edge segment of the polygon x be described by four parameters: horizontal-position, vertical-position, length, and orientation. The components of the transformed polygon x' can be represented similarly. The transformation parameters T can also be represented by four components: change-in-horizontal-position, change-in-vertical-position, scale, and rotation. Let $x = (x, y, l, \theta)$, $x' = (x', y', l', \theta')$, and $T = (\Delta x, \Delta y, s, \Delta\theta)$. Then



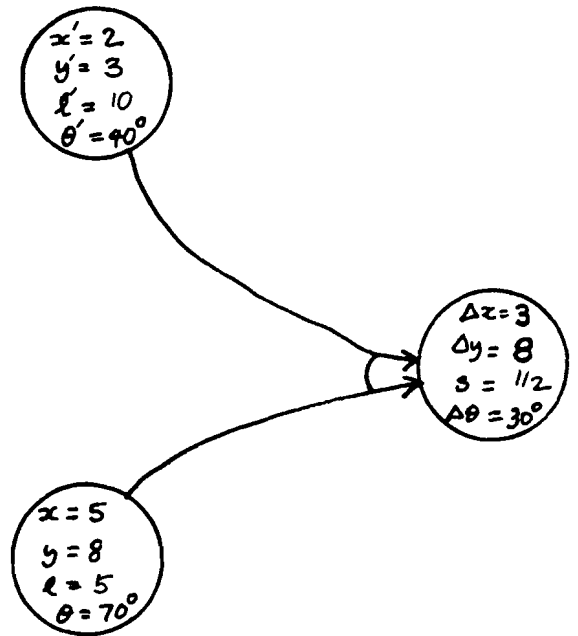
a



b



c



d

Figure 13: (a) Data for a "wrench" displayed as four-parameter (x-location, y-location, length, orientation) vectors. (b) Corresponding data for a simulated image. Here the wrench appears scaled, translated, and rotated together with other background vectors. (c) The principle of detecting a transformation via explicit transformation units. Each (x, x') pair of units is connected to a transformation unit. The unit that receives the most is selected as the appropriate transformation. (d) An example of a specific connection.

$$\Delta\theta = \theta - \theta'$$

$$s = 1/l'$$

$$\Delta x = x - x'.s.\cos\theta - y'.s.\sin\theta$$

$$\Delta y = y + x'.s.\sin\theta - y'.s.\cos\theta$$

These relations can be used to determine connectivity patterns between three networks of four-dimensional value units. As an example of a particular connection, the unit $x' = (2, 3, 10, 40^\circ)$ together with the unit $x = (5, 8, 5, 70^\circ)$ connect to the unit $T = (3, 8, \frac{1}{2}, 30^\circ)$, as shown by Figure 13b. Each pair of units that are on $(s(x) = s(x') = 1)$ provides input for a transformation unit. If there is a subset of (x, x') pairs that are related by the same transformation T , then that T unit will receive the most input.

This strategy was successfully used to recognize two-dimensional shapes in [Ballard, 1981]. However, we were more interested in extending this result to three dimensions. In three dimensions, seven parameters describe the transformation: scale (one), rotation (three), and location (three). Simple calculations show that if there are M units each of different x' , x , and T , there will be $3M^2$ total conjunctive connections and approximately M per unit. In the case where $M = N^k$, where as used earlier, N is the number of just noticeable differences and k is the dimension of the stimulus, the number of units can become prohibitively large. For the three-dimensional case, $(N=100, k=7)$ would require 10^{14} units! Coarse coding can be used to reduce this number and will be especially effective for encoding T units since only one value will be present for most tasks. In general, however, additional economies must be sought.

One general strategy for dealing with high-dimensional parameters is to split them up into subspaces [Ballard, 1984; Ballard and Sabbah, 1983; Hrechanyk and Ballard, 1983]. This is especially attractive in the case of transformations, since explicit algebraic relationships can be found between sub-dimensions, as shown by the earlier equations. (This has recently been done for three-dimensional shapes [Ballard and Tanaka, 1985], but we will stick with two-dimensional shapes in our example, as the graphic displays are more intuitive.) In the two-dimensional example, four-dimensional x , x' and T units can be split into nine subspaces, each of which contains only N^2 units. The image data is represented in terms of three groups of units: (1) position units; (2) orientation and scale units; and (3) coarse resolution position-orientation-scale units. The last group is the group that was costly to represent under the earlier scheme, but in this case the resolution of these units is extremely coarse. This is possible since they work in conjunction with the other fine resolution units. Model units are represented in the same way, with three groups of units like those that represent image data. Interestingly, these groups are similar to cells found in striate cortex. Simple cells are position sensitive, complex cells are orientation sensitive (and, unlike the units in our model, weakly position sensitive), and hypercomplex cells are coarsely position and orientation sensitive. The transformation T is represented by three groups of units: (1) scale-rotation units; (2) rotated position units; and (3) translation units. Figure 14 shows the relationship

between these units. Note that only representative units are drawn. Each such unit is part of a group which collectively covers all values of the particular parameters.

This example reflects computational requirements more than the underlying biology; it is a sufficiency proof that the transform can be computed in parallel using only two-dimensional value units. Nonetheless, since the algebraic relationships are basic, it would not be surprising if they were exploited by the underlying biology. The "answer" is described by the most active rotation and scale unit (5 in Figure 14) and the most active translation unit (9 in Figure 14).

One drawback that the subspaces have is that an image shape that matches correctly in each subspace but does not have the right location/orientation-scale correspondence would be indistinguishable from the correct shape. In other words, ($\uparrow\leftarrow$) would match ($\leftarrow\uparrow$). This is an example of a general phenomena of errors in subspace correspondence known as illusory conjunctions [Treisman and Gelade, 1980]. In our design we implemented constraints that loosely couple the two subspaces in order to reduce the amount of illusory rotation/scale conjunctions. However, this problem is probably solved by selective focus of attention. But since this is a sequential mechanism, it is beyond our present scope. Another helpful constraint is that of rotation scale filtering, which gives more emphasis to length-orientation units that have an associated *on* transform unit and model unit.

Figure 15 shows several of the networks in the computer simulation after one parallel iteration. In this experiment the data of Figure 13c and d was used as the model and retinotopic input, respectively. The model is a caricature of a wrench that appears transformed in the simulated image together with artificially generated noise vectors. Figure 15a and b shows how the wrench looks in terms of the two subspaces (x' : location) and (x'' : rotation and scale), respectively. Figure 15c and d shows that the network correctly computes the answer: single units in each of these spaces receive the most input.

Figure 14.

Figure 15.

Shape recognition is only one of several problems that require the computation of coordinate transformations. Two other examples are sensory motor guidance [Sparks, 1983] and stable perception. The solution for shape perception directly carries over to these other problems, although in each case there will be many other non-transformational issues. For a development of the stable perception issues, see [Feldman, 1982].

5.2 Motion Perception

Almost all computational approaches to vision adopt the use of hierarchies that represent parameters at increasing levels of abstraction. As mentioned in Section 3, there is considerable evidence that the cortical representations use similar hierarchies. These hierarchies were developed independently. The computational approach

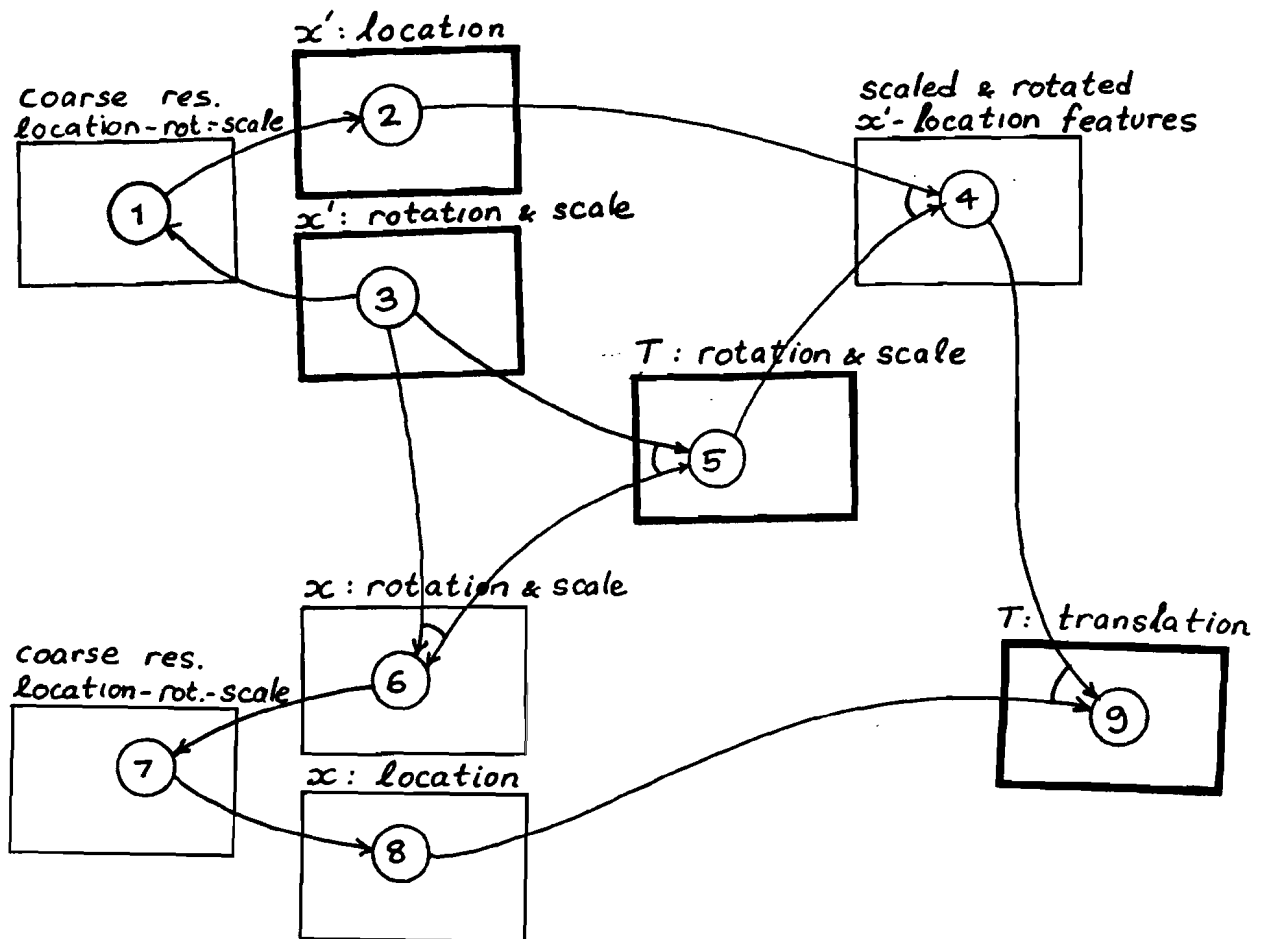



Figure 14: The decomposition of the four-parameter units into nine parameter spaces of two-parameter units. Only representative connections are shown. The boldface notation indicates the value units depicted in a computer simulation in Figure 15.

Input	Output	Type	Connection Rule	Comment
1: (X', Y', L', θ') 3: (l', θ')	2: (x', y') 1: (X', Y', L', θ')	A	$x' = X'$; $y' = Y'$ $L' = l'$, $\theta' = \theta'$	loosely couples two subspaces
2: (x', y'), 5: (s, $\Delta\theta$)	4: (x _r , y _r)	M	$x_r = x'.s.\cos\Delta\theta$ $+ y'.s.\sin\Delta\theta$ $y_r = -x'.s.\sin\Delta\theta$ $+ y'.s.\cos\Delta\theta$	rotate model
4: (x _r , y _r), 8: (x, y)	9: (Δx , Δy)	M	$(\Delta x, \Delta y) = (x_r - x, y_r - y)$	translation matching
3: (l', θ'), 5: (s, $\Delta\theta$)	6: (l, θ)	M	$(l, \theta) = (s.l', \Delta\theta + \theta')$	rotation/ scale filtering
3: (l', θ'), 6: (l, θ)	5: (s, $\Delta\theta$)	M	$(s, \Delta\theta) = (l/l', \theta - \theta')$	rotation/ scale matching
6: (l, θ)	7: (X, Y, L, θ)	A	$l = L$, $\theta = \theta$	loosely couples two subspaces
7: (X, Y, L, θ)	8: (x, y)	A	$X = x$, $Y = y$	


A = additive, connections summed $p_i = \sum w_{ij} s_j$

M = multiplicative, conjunctive connections $p_i = \sum w_{ijk} s_j s_k$


Figure 14 Cont. A table showing the different algebraic relations used to connect the units in the figure.




a



b



c



d

Figure 15: The results of a computer simulation, showing the parameter spaces that were outlined in Figure 14. The input to a unit, p_i , is encoded as a grey level. The data shown in Figure 13a and b were used as the prototype and image, respectively. The objective is to find the correct transformation (scale, rotation) and (x-translation, y-translation) units that relate the prototype to a subset of the image data. The top two pictures show the subspace encoding technique using the prototype as an example. (a) Location units for Figure 13a. (b) Rotation, length units for Figure 13a. The bottom two pictures show the subspace encoding of the transformation space. (c) The correct *rotation-scale* unit receives the most input. (d) The correct *translation* unit receives the most input.

demands an explicit representation and algorithm to compute information necessary to do a specific task. The biological approach is aimed at explaining anatomical and physiological data from the cortical substrate. Yet despite this independence, the hierarchies reveal surprising correspondences. To be specific, we consider the example of motion detection. Table 3 compares computational parameters with plausible biological areas. The biological areas were selected on the basis of single neuron recording data and interconnection hierarchies [Maunsell and Van Essen, 1982].

Table 3: Motion Perception Hierarchy

level	abstraction	computational parameters	plausible biological area
1	image	spatially distributed intensity levels	rods, cones
2	spatial and temporal change	local intensity change due to movement, edges or lighting changes	LGN, V1, V2
3	optic flow and derivatives	physical velocities and spatial derivatives	MT [Allman et al., 1982; Movshon, 1983]
4	rigid motion	parameters of rigid motion: rotation and translation	posterior parietal [Sakata et al., 1980]; superior temporal sulcus [Bruce et al., 1981]

From the computational perspective, each of these levels solves a certain problem: they extract important parameters from the level below and represent them explicitly. The exact choice of representation between levels 1 and 2 is guided by the need to represent both spatial and spatial frequency information [Sakitt and Barlow, 1982] and the need to represent information at different scales [Poggio et al., 1982]. Optic flow is the retinotopic velocity field induced by motion of the observer or objects in the world [Gibson, 1950]. Not all spatio-temporal change is due to movement, e.g., consider changes in illumination strength. Thus optic flow makes explicit the spatio-temporal change that arises from physical movement. The explicit representation of optic flow is an important step, but in order to use motion stimuli fully, the information must be transformed into rigid body parameters. In a value unit model, optic flow vectors turn on units representing global rotation and translation vectors. Information in this form can be used for recognition, eye movements, and navigation.

One could question whether all these levels are really necessary. From pure computational principles, it is possible to bypass levels and compute rigid body motion parameters directly from image intensities; however, the value unit representation scheme demands intermediate levels because relations that are expressed in terms of many parameters involve too many connections. Like the number of units, the number of connections increases exponentially with the number of parameters. Thus in value unit representations, the only scheme which is practical

in terms of numbers of connections must use intermediate representational levels. Keeping the constraints simple by introducing hierarchies has another advantage: there is a better chance that abstract units can evolve naturally during a learning process by being fortuitously connected to highly correlated units representing less abstract information [Barlow, 1983; von der Malsburg and Willshaw, 1977; Hinton and Sejnowski, 1983]. Thus in terms of normal spatio-temporal change, optic flow may be an invariant that is sufficiently probable to be learned by a Hebbian scheme such as [Hinton et al., 1984; Barto et al., 1982], whereas rigid motion parameters may be *insufficiently* correlated. On the other hand, the rigid motion parameters may be highly correlated in terms of optic flow. Once an abstract level has been learned, levels that are even more abstract can be learned.

While the optic flow computation is particularly elegant, similar constraints may be found for rigid body motion parameters. Ballard and Kimball [1983] have shown a formulation that uses temporal variations in the optic flow. In that implementation, rotation is found by assuming that three-dimensional velocity information is available from the time-varying image. Other parametric approaches to representing rigid body motion have used the two-dimensional optic flow image [Prazdny, 1981; O'Rourke, 1981; Ullman and Hildreth, 1983]. In particular situations, such as translatory motion, other interesting parameters can be computed directly from the optic flow [Lee and Reddish, 1981; Hrechanyk and Ballard, 1983]. All these formulations can be translated into the value unit formalism.

Computing Flow

To show how the value unit formalism can handle motion, let us start with the interconnections at the less abstract end of the hierarchy. Specifically, we show how optic flow can be computed from spatio-temporal change. Computing the local velocity at a specific retinal location is done differently depending on whether or not the spatial image intensities are oriented. To consider the two different cases, suppose the image contains an *edge*. In this case, only velocity estimates perpendicular to the edge can be obtained (the well-known "aperture problem"). The other case is when the image contains relatively unoriented intensity patterns. If this image moves, the velocity can be detected by correlation in space and time. However, space-time correlation is costly to implement biologically. A computationally economical solution is to be content with the direction of motion and take advantage of motion blur. To do this one can use oriented pattern detectors provided their integration times are sufficiently long. That is, motion blur produces a spatially-oriented pattern which triggers the detector. These two strategies are depicted in Figure 16. Note that this slow behavior is the opposite to that of the edge detectors which must have fast integration times. Slow integration times in the case of the edge detector will nullify its output.

Table 4: Detector Summary.

τ = unit integration time; V = speed; D = diameter of spatial RF

Type of Value Unit	Integration Time	Spatial Frequency	Oriented
Edge	$\tau \ll V/D$	low	yes
Blur	$\tau \gg V/D$	high	yes
Spot	$\tau \ll V/D$	low	no

In addition to velocity direction, velocity magnitude, or speed, can be detected by linking pairs of spot detectors. The important point here is that speed can be detected with good accuracy given low spatial resolution. These three kinds of value units are described in Table 4. Interestingly, the Edge, Blur, and Spot value units have strong similarities with neurons in the Y-, X-, and W-system as described by [Stone et al., 1979], but these identifications must be considered very speculative at this point.

Up to this point, the temporal response of the value units to the input has not been discussed. For most applications, simply thinking of the response as instantaneous is sufficient to illustrate important behavior. However, the motion example is a case where the temporal behavior of the unit is important. We assume that for most cases the temporal response, expressed as an *integration time* τ , is very fast with respect to the stimulus (as before), but that in the special case of blur units, it is more helpful to have the response time be *slower* than the stimulus. This allows the unit to respond to the blurred photometric pattern.

A circuit that uses these different kinds of value units to compute optic flow is shown in Figure 17. All of the units are spatially indexed, that is, repeated for different retinotopic locations; only the representative units necessary for the computation of one flow vector are shown. Units are on or off according to the model discussed earlier. Specifically, at each point: (1) an edge (θ) unit is on if the input is a linear intensity discontinuity of the appropriate orientation angle; and (2) a blur (θ) unit is on if the input is moving past that location at the appropriate orientation angle [Bandyopadhyay, 1985b]. These are two different ways of signaling velocity direction. Speed is detected by combining coarsely-tuned spot detectors from different spatial locations. Inputs are linked if spatio-temporal summation is required to turn the unit on. The notation " δ " means that the inputs must be different by a fixed time delay δ . If the flow estimate is determined by an edge (θ) unit, additional computation is required to estimate the true flow vector since the velocity estimate is only for the component perpendicular to the edge. One way of combining the information is to use two different estimates. The appropriate connection rule can be worked out simply from a graphical construction [Horn and Schunk, 1981; Movshon, 1983], and Figure 17 shows how to combine two perpendicular components.

Figure 16.

Figure 17.





Figure 16: Two ways to measure optic flow. Top: (a) Moving high frequency patterns indicate flow direction after (b) temporal blurring [Bandyopadhyay, 1985b]. Bottom: (c) Moving edges only indicate flow in the direction perpendicular to the edge. However, two non-parallel edge measurements can be combined, as shown in (d) [Movshon, 1983].

12-11

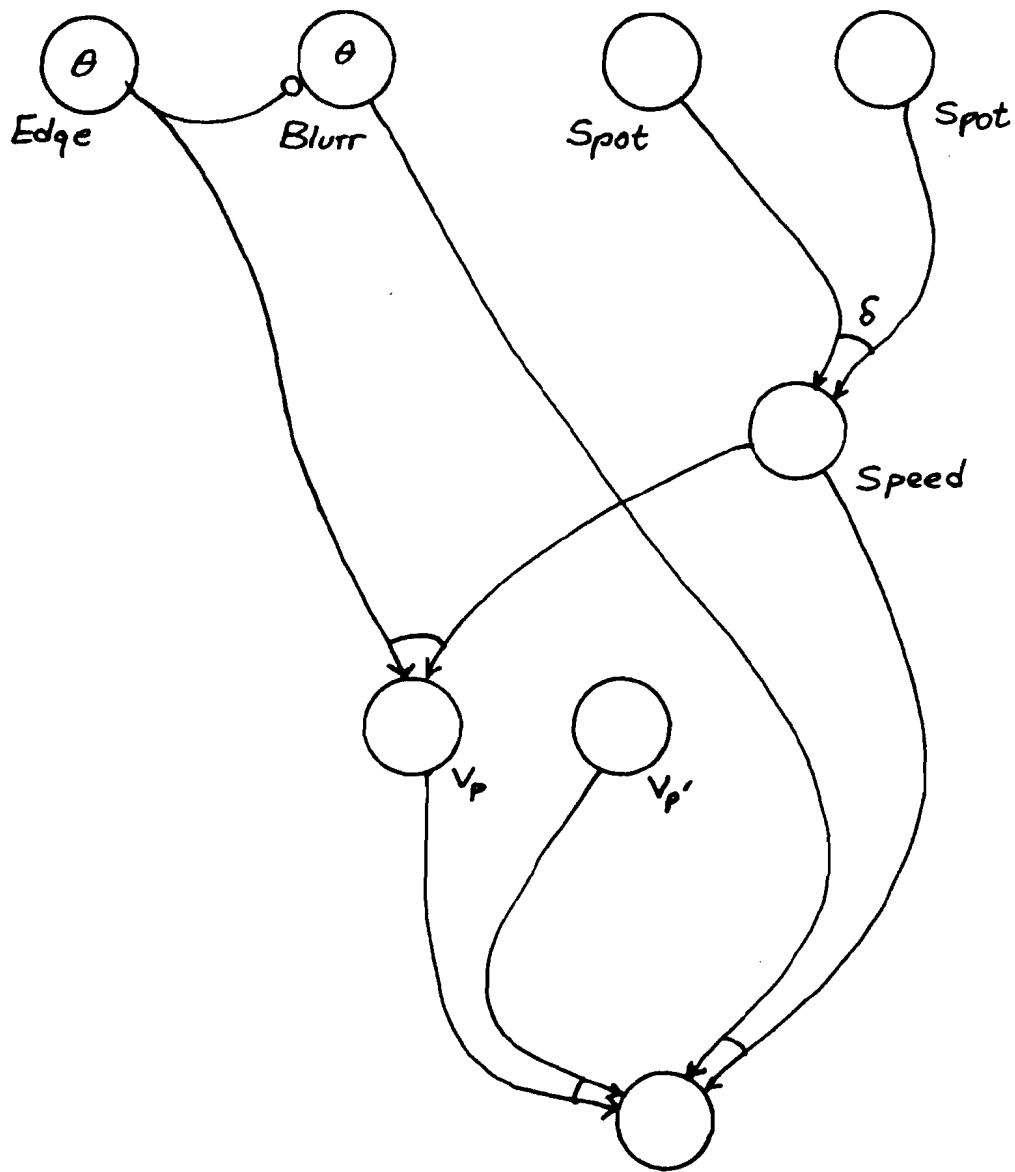


Figure 17: A circuit that implements the different optic flow detection strategies described in the text. Oriented *Edge* units record the perpendicular direction of local motion. These can be combined with a measure of speed to estimate the velocity perpendicular to the edge v_p . Two local estimates, v_p and v_p' , can be combined to obtain an optic flow vector. A separate mechanism uses blur units that correctly measure velocity direction from temporally blurred images. These, when combined with a speed measurement, provide a separate estimate of optic flow.

As shown in Ballard and Kimball [1983], the value-unit approach can be also applied to the more abstract hierarchical levels. One wants to characterize the gestalt of rigid body motion; the solution is to define a parameter space of value units. Functional constraints determine the wiring between value units in the different representational spaces. If the parameter space dimensionality requires an unrealistic number of units, the space is implemented as subspaces of interconnected units. Like the form perception experiments, the motion experiments use rotational and translational subspaces.

In a value unit experiment to detect rigid motion parameters, if depth and optic flow are available, a description of the motion can be created in terms of rigid body parameters. We assumed that related work in stereo [Marr and Poggio, 1976] was successful, in addition to the optic flow networks.

Given flow and depth, three-dimensional (3d) flow can be recovered. From 3d flow, use Newtonian kinematics:

$$v(X,Y,Z) = V_T + \omega x(X-X_c, Y-Y_c, Z-Z_c)$$

where $v(X,Y,Z)$ is the 3d velocity of a point in space, which may be expressed in terms of a rigid translational velocity V_T and a rotation ω about an origin (X_c, Y_c, Z_c) . The summary of our method is as follows:

Step 1: For each pair of velocities $v(X1,Y1,Z1)$ and $v(X2,Y2,Z2)$, build a value unit network to find $\text{dir}(\Omega)$ from $\text{dir}(v(X1,Y1,Z1)) \times \text{dir}(v(X2,Y2,Z2))$

Step 2: Build a value unit network to find $\text{mag}(\Omega)$ and V_T from the above equation, assuming $\text{dir}(\Omega)$ is known

In this description, "dir" and "mag" stand for direction and magnitude, respectively. Step 1 is actually an improvement over the originally proposed method described in [Ballard and Kimball, 1983]. The results of these experiments are shown in Figure 18.


Figure 18.

Different constraints than Ballard and Kimball's [1983] are currently being tested [Prazdny, 1981; O'Rourke, 1981; Ullman and Hildreth, 1983; Lee and Reddish, 1981; Lawton, 1983; Bandyopadhyay, 1985a], but the usefulness of the rigid motion parameters themselves suggests that they will be likely to be represented cortically. Preliminary evidence for neurons sensitive to the appropriate full-field stimuli has been found in the superior temporal sulcus [Bruce et al., 1981] and in posterior parietal cortex [Sakata et al., 1980].


While we have just sketched the principal features of a value unit solution to motion representation, this level of discussion is sufficient to allow us to draw important conclusions. The most important concerns the meaning of the different cortical areas. The larger and larger receptive fields encountered as the anatomic



a



b



c

Figure 18: The next step in the hierarchy uses units to represent subspaces for global rotation and translation parameters. A global parameter is independent of retinotopic position and represents full-field stimuli. The figure shows a result of a computer experiment where three-dimensional velocity vectors provided input to rotation and translation subspaces of units as described in the text. (a) One frame of three-dimensional vectors. (b) A geodesic value unit display representing possible rotation directions, showing that the correct rotation direction unit receives the most input. (c) Three subspaces of units representing the magnitude of rotation combined with the three components of translational velocity. Although many units receive input, the three correct units receive much more input than the others.

...hierarchy is progressed can be explained as a natural consequence of the more abstract parameters that are being represented. The most important point is that hierarchies of units represent visual information of a different nature as implied by Table 3.

As in the form perception example, the motion example shows that physical constraints can be directly represented in terms of the cortico-cortical connections if the constraints only involve a small number of terms. If in fact the constraints used by the brain are multivariate and of high dimensionality, then the value unit model is unrealistic.

To summarize the main points of this section:

- 1) Important problems in visual perception can be readily expressed in connectionist models. Two examples are coordinate mappings and rigid body motion perception. The specific constraints used for each of these cases are constantly being refined, but our claim is that the right constraints will still have value unit implementations.
- 2) The networks that solve these problems make extensive use of hierarchies and parameter subspaces. Both of these techniques help to keep the connection growth problem, introduced by the value unit formalism, under control.

6. Associating Value Units in Different Modalities

The splitting up of large parameter spaces into subspaces is a tremendously important strategy. It makes the value unit model extensible to arbitrary number of dimensions as long as they can be decoupled in some way. On the other hand, it is easy to see that the use of subspaces instead of the high-dimensional space can lead to problems. Consider Figure 19 which represents a subspace of color units and a subspace of form units. The lower group of units differs from the upper in that the lower group are retinotopically indexed whereas the upper group are not. For this reason the upper group are referred to as "de-spaced." Suppose the mechanisms of the previous section (at a much coarser grain) are given a stimulus of a triangular, blue object and a circular, red object. If one considers the subspaces alone as in Figure 19, then the immediate problem is that the associations between the high confidence units have been lost. There is a growing body of psychological evidence that suggests that this is a ubiquitous problem in perception. During fast (~100 ms) tachistoscopic presentations of colored shapes, subjects will make submodality association errors, so-called "illusory conjunctions," but given more time they can usually correct their mistakes [Treisman and Gelade, 1980]. The illusory conjunction experiments suggest both that the separate groups of value units is a real problem and that given additional processing time beyond a few hundred milliseconds, there is a solution to this problem. Furthermore, the psychological data suggests that the conjunctive mechanism is sequential. In experiments with displays of colored letters, subjects were asked questions of the form "Is there a red triangle?" In these experiments, the processing time increases linearly with the number of visual tokens. However, if the question is "Is there a triangle?" the processing time is independent of the number of tokens. One value unit solution to this problem, suggested by Feldman and Ballard [1982], uses spatial focus units to couple the different submodalities of color and form.

Figure 19.

In accordance with the psychological data, this solution assumes that the overall computational behavior of the networks is *sequential* within the submodalities. That is, priming a despaced unit, say "triangular," causes increased input to the appropriate spatial focus units which then: (a) inhibit other spatial locations; and (b) help computations at the spatial focus locations. Thus processing can be restricted to the appropriate spatial locus. We hasten to say that this explanation captures only the gross features of Treisman's data: the fact of illusory conjunctions and sequential processing. Additional circuitry is necessary to simulate all the precise timing effects. Crick [1984] has suggested a neurobiological locus for this mechanism that implicates thalamic nuclei as well as cortical areas.

The ability to associate high-confidence value units provides a mechanism necessary to isolate a particular stimulus. This is especially important in the value unit model for the following reason. Consider the example of picking up or looking at an object. Since we postulate that all the cortex is represented in value units, this control must be achieved through hard-wired connections from sensory to motor cortex. If the connections involve intermediate units, they are also value units. This means that if the stimulus is not isolated as a single, localized group of value units,




Figure 19: Using spatial units to couple de-spaced color and motion units. By selecting a specific modality, for example color, its corresponding spatially-coincident feature may be found even if it is not *directly* coupled. In this case, non-spatially indexed color units excite spatially indexed color units which in turn are coupled to spatially indexed shape units. This preferentially excites the non-spatially indexed shape units that correspond to the color. The enhanced links are shown in boldface in the figure.

the motor system will be compromised by receiving parallel inconsistent inputs. An example would be a frog snapping at the "average fly" produced by two concurrent moving spot stimuli [Didday, 1976]. Although this did occur in Didday's detailed simulation, it was rare; interestingly, he used network-like control structure to isolate a single stimulus.

Maintaining Associations

Although for some situations it is sufficient to isolate a single stimulus, more complex situations require the ability to:

- 1) use two or more associations at the same time; or
- 2) remember previous associations.

For these kinds of problems a number of solutions have been proposed: (a) shared links [Feldman, 1982]; (b) dedicated links [Feldman, 1982]; (c) synchronously firing units [Abeles, 1982; Crick, 1984]; and (d) fast, short-term synaptic weight changes [von der Malsburg, 1981]. Hinton and Sejnowski's weight change formalism [1983] is also a candidate for (d), but at present has been used to model long-term changes. This paper focuses on fast, short-term weight changes. A weight change is fast if it takes place within two to fifty milliseconds. Such weight changes provide a passive memory for value unit networks that avoids conjunctions between previous and ongoing computations.

The idea of fast, short-term synaptic weight changes has had a history of varying popularity. The novel addition to the idea that we propose is that the change in weights cannot be random or correlated with any repeated firings but instead must be more structured. The way we propose to do this is to use convergence states as a trigger for such modifications. The difficulty with an unstructured weight modification scheme is that it may interfere with the base relaxation computations which require the weights to be kept constant. This requirement follows from the fact that during the short convergence cycles of the relaxation, inappropriate units may temporarily have high confidence values. An example is the case of a local visual feature that turns out to be inconsistent with more global context. A certain number of cycles are required to allow more global value units to "overrule" this local evidence [Sabbah, 1981; 1982]. If weights were allowed to change very quickly, then the local evidence would continue to predominate. Thus the main point is that the weight modifications must not occur during short-term convergence, as the intermediate and final states may be different.

The other requirements are: first, that the weight change must be rapid relative to a few hundred milliseconds, and second, that the weight change is a state change that should persist for a substantial period (minutes) while other computations are carried out. The first requirement is strictly for performance; other computations which depend on the weight change are held up until it occurs. The second requirement characterizes the weight change as a system state change that should persist long enough to be useful in several other computations.

To illustrate the problem and solution further, we develop the color-form example. Consider again the case of four objects which are described by the different combinations of two shapes and two colors. A connectionist method of representation would allocate a unit for each of the objects, as shown in Figure 20. Conjunctive connections can be used for each of the stimulus pairs. To change the weights, one can first use the focusing method described earlier to isolate an individual pairing, and only at that point: (a) increase the weights of all the active bindings; and (b) decrease the weights of all the inactive bindings. After this has been done, only the correct objects will be activated even though all the combinations of features are active. (In this example, although the symbols on the units are the same for the purposes of discussion, the actual information will be different at each of the three levels. As in Figure 19, the first two units represent retinotopically indexed information and de-spaced information, respectively. The third level has more coarsely coded information appropriate for object classes. For example, all canonical triangles at level 2 will map into the triangle category node at level 3. If categories were not used at level 3, one would quickly run out of units.)

Figure 20.

Special-purpose connection systems required to implement weight changes may reside in the hippocampus, but the evidence is also still very speculative. The hippocampus involves multimodal connections from almost all cortical areas. Furthermore, it is one of the few places where long-term synaptic potentiation has been observed. Also, interfering with the hippocampus impairs the ability to do tasks which require converting short-term associations to longer-term memories. A host of possibilities for a chemical mechanism are available--most likely calcium channels or peptides.

To summarize the main points of this section:

- 1) The use of subspaces is ubiquitous and raises the difficult problem of how to associate units in different subspaces that are part of the same percept. Several connectionist mechanisms which solve this problem are under study. The possibilities emphasized in this paper are: (a) sequential focusing to isolate associations; and (b) fast synaptic weight changes to maintain associations.
- 2) Fast synaptic weight changes are a passive form of memory that allows the same network to be used for additional computation. In active forms of memory the separation of the remembered state from the current state is more difficult to control.




Figure 20: The basis for the cross-talk problem. The object network must be connected to handle feature cross products and yet keep individual bindings distinct. *Sequentially* isolating the bindings allows the weights on the appropriate pairings to be changed to reflect the spatial correspondences. The changed weights are indicated by boldface links. The weights are reset when the bindings are no longer appropriate. The higher the level, the more abstract the encoding. Level 1 encodes retinotopic information. Level 2 encodes de-spaced object-centered information. Level 3 encodes categorical information.

7. Conclusion

We have argued that the value unit model is compelling in light of current neurological findings and computational studies, but what are the alternatives to value unit connectionism? To address this question, let us briefly summarize the features of the value unit design:

- 1) It is a representation that allows parallel access to different values of a variable. This allows parallel computation within the observed 300 ms human responses.
- 2) It is a representation that is independent of particular constraints. Although our examples came from vision, value units could be a lingua franca for the cortex that allows easy communication across different modalities. Value units also remove the distinction between symbolic and numerical data. Success in language modeling by [McClelland and Rumelhart, 1981] and the general features of somatosensory cortex [Woolsey, 1981] suggest this ubiquity.
- 3) It is compatible with a computational model that tries to minimize energy in satisfying constraints. This model, developed out of work by several groups (Hopfield, Kirkpatrick et al., Hinton and Sejnowski, Geman and Geman), has the advantage that it divorces the representation from the computation. Thus these two issues can be addressed separately.
- 4) It is an encoding that emphasizes collections of small parameter descriptions. The decoupling of the vision problem into local constraints that can be represented in terms of low-dimensional variables has led to major advances [Maloney, 1984; Pentland, 1984]. Thus value units seem to be a representation that allows hard problems to be encoded into computationally tractable forms.
- 5) Hierarchies of value unit encodings are robust; networks still compute even when large portions are removed.
- 6) Hierarchies of value unit encodings can exhibit enormous diversity. One criticism that might be leveled is that 10^{11} neurons would not be sufficient to encode the necessary experiences. However, this is where hierarchies help. Crudely put, hierarchies of value units form a kind of "numeric" representation, with abstract units forming "higher-order bits" and less abstract units forming "lower-order bits" that can be composed. In a similar way, subspaces also expand the number of possibilities by allowing compositions. Thus the number of objects representable would be: (the number of object token units) \times (the number of color units) \times (the number of shape units) \times (the number of motion units), etc.

- 7) Hierarchies of value units that represent specific invariants that can be computed in parallel represent a straightforward answer to the Gestaltists: the things that naturally organize are just those that have explicit small-parameter descriptions.

Given these advantages, let us consider some alternatives. One alternative is a Von-Neumann-like design, but we can reject this as being too slow when built in neural circuitry. Parallel Von Neumann computers are also infeasible, owing to the difficulty in satisfying property (2). One of the big current problems with networks of Von Neumann machines is the time taken in interpreting different message protocols.

A second alternative is variable encoding. Analog computers are designed on this principle, and there is no reason why differential equations could not be directly encoded in neural circuitry. Current models of the eye movement control circuitry are based on analog servomechanisms and have made several important predictions [Robinson, 1978]. The advantage of the analog encoding is that it is more compact than the value unit encoding (but probably only by about a factor of 10^2 , since the firing rate is extremely band-limited). Variable units have two important disadvantages: (a) only one value of a variable can be accessed at a time; and (b) the circuitry tends to be more delicate (adding or removing pieces effects performance unpredictably). Nonetheless, just as the thalamus uses some value encoding, the cortex could use some variable encoding. However, most single cell electrophysiological recording data would rule out variable units.

One controversial aspect of the value unit hypothesis is the encoding of values in a few units. This has become known as the "localist" hypothesis, as opposed to the "distributed" hypothesis that suggests that encodings involve many hundreds or more units. The N^k argument sharpens this debate. Naturally, many hundreds of units will be involved in a percept; the crucial question is: how are small groups of parameters handled? Keeping the number of units that represent a parameter vector small will facilitate parallel computations, since nearby values do not interfere, and may simplify the connection problem, since another network that requires the value need only connect to a few units. Besides these advantages, single cell recording data seems to be in favor of a localist encoding. It certainly is for the highly retinotopic areas such as V1 and V2. We argue that once the correct parameters are identified it will extend to other extra-striate areas as well.

What are the advantages of locality? In connectionist terms, to implement the useful relationships among sensory and motor parameters, it is enormously useful to have value units that are *similar in value* be *physically proximal*. However, as the value units become increasingly abstract, the notion of value locality becomes more vague and the physical locality less imperative. Extremely abstract units may form a diffuse network that is scattered throughout the cortex and obeys no regular pattern. In this case the primary-secondary indexing concept may not be useful, but we believe the concepts of value units and topological locality could still prove useful. Skimpy evidence comes from patients with lesions of the corpus callosum who exhibit very discrete functional losses (e.g., [Dimond et al., 1977]).

To conclude on a less technical note, for a long time strictly computational models have not had a significant interaction with basic studies in the neurosciences. However, new discoveries in both these areas are leading to a renaissance of attempts to bridge these disciplines. We hope that this paper will spark additional interest in interdisciplinary studies.

8. Acknowledgements

This paper was an outgrowth of a University of Rochester Bridging Fellowship which allowed the author to spend one year studying in the Departments of Anatomy and Physiology with Paul Coleman. During that year many people helped initiate the effort that culminated in the paper. In particular, Paul critiqued many earlier drafts and the paper, and substantially improved the presentation. The author also wishes to thank Mike King, Joanne Albano, Dorothy Flood, Chris Block, Gloria Hoffman, Tanya Pasternak, and the Center for Visual Science seminar group, particularly Peter Lennie, Dave Williams, and Walt Makous. Also, Mike, Joanne, and Dorothy extensively reviewed earlier drafts, as did Francis Crick of the Salk Institute. In addition, on the computer science side, Jerry Feldman and Chris Brown provided helpful critique. The author wishes to acknowledge partial support from the National Science Foundation (MCS 820320) and the National Institutes of Health (HL21253). Finally, many thanks go to Peggy Meeker for preparing the numerous drafts of the manuscript.

9. References

- Abeles, M. *Local Cortical Circuits: An Electrophysiological Study*. New York: Springer-Verlag, 1982.
- Allman, J.M., J.F. Baker, W.T. Newsome, and S.E. Petersen, "Visual topography and function." In C.N. Woolsey (Ed). *Cortical Sensory Organization: Volume 2*. Clifton, NJ: Humana Press, 1982.
- Ballard, D.H., "Generalizing the Hough transform to detect arbitrary shapes," TR 55, Computer Science Dept., U. Rochester, 1979; *Pattern Recognition 13*, 2, April 1981.
- Ballard, D.H., "Parameter networks: Towards a theory of low-level vision," TR 75, Computer Science Dept., U. Rochester, 1981; *Proc, 7th IJCAI*, Vancouver, B.C., August 1981; *Artificial Intelligence 22*, 235-267, 1984.
- Ballard, D.H. and P.J. Hayes, "Parallel logical inference," to appear, *Proc.*, Cognitive Science Society Conference, Boulder, CO, June 1984.
- Ballard, D.H. and O.A. Kimball, "Rigid body motion from depth and optical flow," TR 70, Computer Science Dept., U. Rochester, November 1981; *CGIP Special Issue on Computer Vision*, 1983.
- Ballard, D.H. and D. Sabbah, "On shapes," *Proc.*, 7th IJCAI, Vancouver, B.C., August 1981; to appear, *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 1983.
- Ballard, D.H. and H. Tanaka, "Frame-based form perception: Constraints, algorithms, implementation," submitted to the *9th Int'l. Joint Conference on Artificial Intelligence*, Los Angeles, CA, August 1985.
- Ballard, D.H., G.E. Hinton, and T.J. Sejnowski, "Parallel visual computation," *Nature 306*, 21-26, 3 November 1983.
- Bandyopadhyay, A., "A connectionist model of rigid body motion perception," forthcoming TR, Computer Science Dept., U. Rochester, 1985a.
- Bandyopadhyay, A., "Spatio temporal blur paths: A representation for image motion," submitted to Conference on Computer Vision and Pattern Recognition, 1985b.
- Barlow, H.B., "Single units and sensation: A neuron doctrine for perceptual psychology?" *Perception 1*, 371-394, 1972.
- Barlow, H.B., "Critical limiting factors in the design of the eye and visual cortex," *Proc. R. Soc. Lond. B212*, 1-34, 1981.
- Barlow, H.B., personal communication, 1983.
- Barto, A.G., R.S Sutton, and C.W. Anderson, "Adaptive neuron-like elements that can solve difficult learning control problems," COINS TR 82-20, U. Massachusetts, 1982.
- Blasdel, G.G., D. Fitzpatrick, and J.S. Lund, "Organization and intracortical connectivity of layer IV in macaque striate cortex," *Proc.*, 13th Annual Meeting, Society for Neuroscience, Boston, MA, November 1983.

- Brady, M., "Computational approaches to image understanding," AI Memo 653, AI Lab, MIT, October 1981; *Computing Surveys*, 1982.
- Brodmann, K. *Vergleichende Lokalisationslehre der Grosshirnrinde in ihren Prinzipien dargestellt auf Grund des Zellenbaues*. Leipzig: J.A. Barth, 1909.
- Bruce, C., R. Desimone, and C.G. Gross, "Visual properties of neurons in a polysensory area in superior temporal sulcus of the macaque," *J Neurophysiology* 46, 2, August 1981.
- Burton, H. and C.J. Robinson, "Organization of the S II parietal cortex: Multiple somatic sensory representations within and near the second somatic sensory area of cynomolgus monkeys," in C.N. Woolsey (Ed). *Cortical Sensory Organization Volume 1: Multiple Somatic Areas*. Clifton, NJ: Humana Press, 1981.
- Cajal, S. Ramon y. *Histologie du Systeme Nerveux de l'Homme et des Vertebres*. Paris: Maloine, 1911.
- Cowey, A., "Why are there so many visual areas?" in F.O. Schmitt, F.G. Worden, G. Adelman, and S.G. Dennis (Eds). *The Organization of the Cerebral Cortex*. Cambridge, MA: The MIT Press, 1981.
- Crick, F., "The function of the thalamic reticular complex: The searchlight hypothesis," *Proc. National Academy of Sciences USA* 81, 4586-4590, 1984.
- Curcio, C.A. and J.K. Harting, "Organization of pulvinar afferents to area 18 in the squirrel monkey: Evidence for stripes," *Brain Research* 143, 155-161, 1978.
- Cynader, M.S., J. Matsubara, and N.V. Swindale, "Surface organization of functional and topographic maps in cat visual cortex," *Proc.*, 13th Annual Meeting, Society for Neuroscience, Boston, MA, November 1983.
- De Valois, K.K., *Vision Res.* 17, 1057, 1977.
- Didday, R.L., "A model of visuomotor mechanisms in the frog optic tectum," *Math. Biosci.* 30, 169-180, 1976.
- Dimond, S.J., R.E. Scammell, E.Y.M. Brouwers, and R. Weeks, "Functions of the centre section (trunk) of the corpus callosum in man," *Brain* 100, 543-562, 1977.
- Dow, B.M. and R. Bauer, "Retinotopy and orientation columns in the monkey: A new model," *Proc.*, 13th Annual Meeting, Society for Neuroscience, Boston, MA, November 1983.
- Eccles, J.C. *The Physiology of Nerve Cells*. Baltimore: Johns Hopkins U. Press, 1957.
- Feldman, J.A., "Four frames suffice: A provisional model of vision and space," TR 99, Computer Science Dept., U. Rochester, September 1982.
- Feldman, J.A., "Memory and change in connection networks," TR 96, Computer Science Dept., U. Rochester, October 1981.
- Feldman, J.A. and D.H. Ballard, "Connectionist models and their properties," *Cognitive Science* 6, 205-254, 1982.
- Freuder, E.C., "Synthesizing constraint expressions," *CACM* 21, 11, 958-965, November 1978.

- Geman, S. and D. Geman, "Stochastic relaxation, Gibbs distributions, and the Bayesian restoration of images," TR, Brown U., September 1983; also appeared in *IEEE Trans. PAMI*, January 1985.
- Gibson, J.J. *The Perception of the Visual World*. Boston: Houghton-Mifflin, 1950.
- Gilbert, C.D., presentation at Rochester Neuroscience Conference, 1982.
- Gilbert, C.D. and T.N. Wiesel, "Clustered intrinsic connections in cat visual cortex," *J. Neurosciences* 3, 5, 1116-1133, May 1983.
- Goldberg, M., personal communication, 1982.
- Goldman-Rakic, P.S. and M.L. Schwartz, "Interdigitation of contralateral and ipsilateral columnar projections to frontal association cortex in primates," *Science* 216, 755-757, 14 May 1982.
- Golgi, C., "Di una nuova reasione apparentemente nera dell cellule nervose cerebrali ottenuta col bichloruro di mercurio," *Arch. Sci. Med.* 3, 1-7, 1879.
- Gross, C.G., C.J. Bruce, R. Desimone, J. Fleming, and R. Gattass, "Cortical visual areas of the temporal lobe: Three areas in the macaque." In C.N. Woolsey (Ed). *Cortical Sensory Organization: Volume 2*. Clifton, NJ: Humana Press, 1982.
- Hebb, D.O. *The Organization of Behavior*. New York: Wiley, 1949.
- Hinton, G.E., "Shape representation in parallel systems," *Proc.*, 7th IJCAI, 1088-1096, Vancouver, B.C., August 1981.
- Hinton, G.E. and T.J. Sejnowski, "Optimal perceptual inference," *Proc.*, IEEE Computer Vision and Pattern Recognition Conf., 448-453, Washington, DC, June 1983.
- Hinton, G.E., T.J. Sejnowski, and D.H. Ackley, "Boltzmann machines: Constraint satisfaction networks that learn," TR CMU-CS-84-119, Computer Science Dept., Carnegie-Mellon U., May 1984.
- Hopfield, J.J., "Neural networks and physical systems with emergent collective computational abilities," *Proc., Nat'l. Acad. Sciences USA* 79, 2554-2558, 1982.
- Horn, B.K.P and B.G. Schunck, "Determining optical flow," AI Memo 572, AI Lab, MIT, April 1980; *Artificial Intelligence* 17, 1-3, 185-204, August 1981.
- Hrechanyk, L.M. and Ballard, D.H., "A connectionist model for shape perception," Computer Vision Workshop, Ringe, NH, August 1982; also appeared as "Viewframes: A connectionist model of form perception," *Proc.*, DARPA Image Understanding Workshop, Arlington, VA, June 1983.
- Hubel, D.H. and T.N. Wiesel, "Receptive fields, binocular interaction and functional architecture in the cat's visual cortex," *J Physiol (Lond.)* 160, 106-154, 1962.
- Hubel, D.H. and T.N. Wiesel, "Shape and arrangement of columns in cat's striate cortex," *J Physiol. (Lond.)* 165, 559-568, 1963.
- Hubel, D.H., T.N. Wiesel, and M.P. Stryker, *J. Comp. Neurol.* 177, 361, 1978.
- Hummel, R. and S. Zucker, "On the foundations of relaxation labeling processes," TR 80-7, McGill U., Montreal, 1980; *IEEE Trans. Pattern Analysis and Machine Intelligence*, 1983.

- Juliano, S.L., O. Favorov, and B.L. Whitsel, "Reproducibility of 2DG patterns in monkey SI and their relationship to single unit mapping data," *Proc.*, 13th Annual Meeting, Society for Neuroscience, Boston, MA, November 1983.
- Kaas, J.H., R.J. Nelson, M. Sur, and M.M. Merzenich, "Organization of somatosensory cortex in primates," in F.O. Schmitt, F.G. Worden, G. Adelman, and S.G. Dennis (Eds). *The Organization of the Cerebral Cortex*. Cambridge, MA: The MIT Press, 1981.
- Kemperman, J., "Recovering multidimensional punctate data from projections," unpublished ms., Mathematics Dept., U. Rochester, 1982.
- Kirkpatrick, S., C.D. Gelatt, and M.P. Vecchi, "Optimization by simulated annealing," *Science* 220, 671-680, 1983.
- Lawton, D.T., "Processing restricted sensor motion," *Proc.*, DARPA Image Understanding Workshop, Arlington, VA, June 1983.
- Lee, D.N. and P.E. Reddish, "Plummeting gannets: A paradigm of ecological optics," *Nature* 293, 24 September 1981.
- Livingstone, M.S. and D.H. Hubel, "Anatomy and physiology of a color system in the primary visual cortex," *J. Neurosciences* 4, 1, 305-356, January 1984.
- Lund, J.S., "Intrinsic organization of the primate visual cortex, area 17, as seen in Golgi preparations," in F.O. Schmitt, F.G. Worden, G. Adelman, and S.G. Dennis (Eds). *The Organization of the Cerebral Cortex*. Cambridge, MA: The MIT Press, 1981.
- Maloney, L., Ph.D. dissertation, Dept. of Psychology, Stanford U., December 1984.
- Marr, D. *Vision*. H.W. Freeman, 1982.
- Marr, D., "Representing visual information," in A.R. Hanson and E.M. Riseman (Eds). *Computer Vision Systems*. NY: Academic Press, 1978.
- Marr, D. and T. Poggio, "From understanding computation to understanding neural circuitry," *Neuroscience Research Progress Bulletin* 15, 470-488, 1976.
- Maunsell, J.H.R. and D.C. Van Essen, "The connections of the middle temporal visual area (MT) and their relationship to a cortical hierarchy in the macaque monkey," Div. of Biology, California Institute of Technology; submitted to *J. Neurosciences*, 1982.
- McClelland, J.L. and D.E. Rumelhart, "An interactive activation model of the effect of context in perception: Part 1," *Psychological Review*, 1981.
- McCulloch, W.S. and W. Pitts, "A logical calculus of the ideas immanent in nervous activity," *Bulletin of Mathematical Biophysics* 5, 115-137, 1943.
- Mishkin, M., L.G. Ungerleider, and K.A. Macko, "Object vision and spatial vision: Two cortical pathways," *Trends in NeuroSciences*, October 1983.
- Mitchison, G. and F. Crick, "Long axons within the striate cortex: Their distribution, orientation, and patterns of connection," *Proc. Nat'l. Acad. Sci. USA* 79, 3661-3665, June 1982.
- Montero, V.M., "Topography of the cortico-cortical connections from the striate cortex in the cat," *Brain, Behav Evol* 18, 194-218, 1981.

- Mountcastle, V.B., "An organizing principle for cerebral function: The unit module and the distributed system," in G.M. Edelman and V.B. Mountcastle (Eds). *The Mindful Brain*. Massachusetts: MIT Press, 1978.
- Movshon, J.A., "Analysis of visual motion," *Proc.*, Conference on Vision, Brain, and Cooperative Computation, U. Massachusetts, May 1983.
- Olberg, Robert M., "Object- and self-movement detectors in the ventral nerve cord of the dragonfly," *J Comparative Physiology* 141, 327-334, 1981a.
- Olberg, Robert M., "Parallel encoding of direction of wind, head, abdomen, and visual pattern movement by single interneurons in the dragonfly," *J Comparative Physiology* 142, 27-41, 1981b.
- O'Rourke, J., "Dynamically quantized spaces for focusing the Hough transform," *Proc.*, 7th IJCAI, 737-739, Vancouver, B.C., August 1981.
- Pentland, A.P., "Shading into texture," *Proc.*, *Nat'l. Conf. on Artificial Intelligence (AAAI-84)*, 269-273, Austin, TX, August, 1984.
- Prazdny, K., "A simple method for recovering relative depth map in the case of a translating sensor," *Proc.*, 7th IJCAI, 698-699, Vancouver, B.C., August 1981.
- Pasternak, T., W.H. Merigan, and J.A. Movshon, "Motion mechanisms in strobe-reared cats: Psychophysical and electrophysical measures," *Acta Psychologica* 48, 321-332, 1981.
- Poggio, T., H.K. Nishihara, and K.R.K. Nielsen, "Zero-crossings and spatiotemporal interpolation in vision: Aliasing and electrical coupling between sensors," AI Memo 675, MIT, May 1982.
- Prager, J.M., "Extracting and labeling boundary segments in natural scenes," *IEEE Trans. PAMI* 2, 1, 16-27, January 1980.
- Rakic, P., "Developmental events leading to laminar and areal organization of the neocortex," in F.O. Schmitt, F.G. Worden, G. Adelman, and S.G. Dennis (Eds). *The Organization of the Cerebral Cortex*. Cambridge, MA: The MIT Press, 1981.
- Rakic, P., "Intrinsic and extrinsic factors influencing the shape of neurons and their assembly into neuronal circuits," in P. Seeman and G.M. Brown (Eds). *Frontiers in Neurology and Neuroscience Research*. Toronto: U. Toronto Press, 1974.
- Robinson, D.A., "The functional behavior of the peripheral oculomotor apparatus: A review," in G. Kommerell (Ed). *Disorders of Ocular Motility: Neurophysiological and Clinical Aspects*. Munich: J.F. Bergman, 1978.
- Rockland, K.S. and J.S. Lund, "Widespread periodic intrinsic connections in the tree shrew visual cortex," *Science* 215, 1532-1534, 19 March 1982.
- Rosenblatt, F., "The perceptron: A probabilistic model for information storage and organization in the brain," *Psychological Review* 65, 386-407, November 1958.
- Rosenfeld, A., R.A. Hummel, and S.W. Zucker, "Scene labelling by relaxation operations," *IEEE Trans. SMC* 6, 420, 1976.
- Sabbah, D., "Design of a highly parallel visual recognition system," 7th IJCAI, Vancouver, B.C., Canada, 1981.

- Sabbah, D., "A connectionist approach to visual recognition," TR 107 and Ph.D. thesis, Computer Science Dept., U. Rochester, April 1982.
- Sakata, H., H. Shibutani, and K. Kawano, "Spatial properties of visual fixation neurons in posterior parietal association cortex of the monkey," *J. Neurophysiology* 43, 6, 1654-1672, June 1980.
- Sakitt, B. and H. Barlow, "A model for the economical encoding of the visual image in cerebral cortex," *Biol. Cybern.* 43, 97-108, 1982.
- Shaw, G.L., E. Harth, and A.B. Scheibel, "Cooperativity in brain function: Assemblies of approximately 30 neurons," *Exp. Neurology*, 1982.
- Shaw, G., P. Renaldi, and J. Pearson, *Exptl. Neurology* 79, 293-298, 1983.
- Shepherd et al., *Proc., Nat'l. Acad. Sci. USA*, to appear, 1985.
- Singer, W., "Topographic organization of orientation columns in the cat visual cortex: A deoxyglucose study," *Experimental Brain Research* 44, 431-436, 1981.
- Sparks, D., "The role of the primate superior colliculus in sensorimotor integration," *Proc., Conference on Vision, Brain, and Cooperative Computation*, U. Massachusetts, May 1983.
- Stone, J., B. Dreher, and A. Leventhal, "Hierarchical and parallel mechanisms in the organisation of the visual cortex," *Brain Res. Reviews* 1, 345-394, 1979.
- Szentagothai, J., "The local neuronal apparatus of the cerebral cortex," in P.A. Buser and A. Roguel-Buser (Eds). *Cerebral Correlates of Conscious Experience*. Amsterdam: North Holland Publishing Co., 1978a.
- Szentagothai, J., "Specificity versus (quasi-) randomness in cortical connectivity," in M.A.B. Brazier and H. Peutsch (Eds). *Architectonics of the Cerebral Cortex* (Int'l. Brain Res. Organization Mon. Series, Vol. 3). New York: Raven Press, 1978b.
- Tootell, R.B.H., M.S. Silverman, R.L. DeValois, and G.H. Jacobs, "Functional organization of the second cortical visual area (V2) in the primate," *Science* 220, 737-739, 1983.
- Treisman, A.M. and G. Gelade, "A feature-integration theory of attention," *Cognitive Psychology* 12, 97-136, 1980.
- Tusa, R.J. and L.A. Palmer, "Retinotopic organization of areas 20 and 21 in the cat," *J Comparative Neurology* 193, 147-164, 1980.
- Tusa, R.J., A.C. Rosenquist, and L.A. Palmer, "Retinotopic organization of areas 18 and 19 in the cat," *J Comparative Neurology* 185, 4, June 1979.
- Ullman, S., "Relaxation and constrained optimization by local processes," *CGIP* 10, 115-125, 1979.
- Ullman, S. and E. Hildreth, "The measurement of visual motion," in O.J. Braddick and A.C. Sleigh (Eds). *Physical and Biological Processing of Images*. London: Springer-Verlag, 1983.
- Van Essen, D.C. and J.H.R. Maunsell, "Hierarchical organization and functional streams in the visual cortex," *Trends in Neurosciences* 6, 9, 370-375, September 1983.

- Van Essen, D.C., W.T. Newsome, and J.L. Bixby, "The pattern of interhemispheric connections and its relationship to extrastriate visual areas in the macaque monkey," *J. Neuroscience* 2, 265-283, 1982.
- von der Malsburg, Internal Report 81-2, Dept. of Neurobiology, Max Planck Institut for Biophysical Chemistry, Goettingen, 1981.
- von der Malsburg, C. and D.J. Willshaw, "How to label nerve cells so that they can interconnect in an ordered fashion," *Proc., National Academy of Science USA*, 74, 11, 5176-5178, November 1977.
- Weller, R.E. and J.H. Kaas, "Cortical and subcortical connections of the visual cortex in primates." In C.N. Woolsey (Ed). *Cortical Sensory Organization: Volume 2*. Clifton, NJ: Humana Press, 1982.
- Woolsey, C.N. (Ed). *Multiple Somatic Areas (Cortical Sensory Organization, Vol. 1)*. Clifton, NJ: Humana Press, 1981.
- Zeki, S.M., "Uniformity and diversity of structure and function in rhesus monkey prestriate visual cortex," *J Physiology* 277, 273-290, 1978.

