LETTER Communicated by Klaus Obermayer

Feedback Decoding of Spatially Structured Population Activity in Cortical Maps

Nicholas V. Swindale
swindale@interchange.ubc.ca
Department of Ophthalmology and Visual Sciences, University of British Columbia, Vancouver, BC, Canada, V5Z 3N9

A mechanism is proposed by which feedback pathways model spatial patterns of feedforward activity in cortical maps. The mechanism can be viewed equivalently as readout of a content-addressable memory or as decoding of a population code. The model is based on the evidence that cortical receptive fields can often be described as a separable product of functions along several dimensions, each represented in a spatially ordered map. Given this, it is shown that for an *N*-dimensional map, accurate modeling and decoding of x^N feedforward activity patterns can be done with *Nx* fibers, *N* of which must be active at any one time. The proposed mechanism explains several known properties of the cortex and pyramidal neurons: (1) the integration of signals by dendrites with a narrow tangential distribution, that is, apical dendrites; (2) the presence of fast-conducting feedback projections with broad tangential distributions; (3) the multiplicative effects of attention on receptive field profiles; and (4) the existence of multiplicative interactions between subthreshold feedforward inputs to basal dendrites and inputs to apical dendrites.

1 Introduction

Activity in the visual cortex, and probably other cortical areas, represents, through sets of overlaid maps, varied combinations of stimulus features, such as orientation, position in visual field, direction of motion, and spatial frequency (Hubel & Wiesel, 1977; Shmuel & Grinvald, 1996; Weliky, Bosking, & Fitzpatrick, 1996; Hübener, Shoham, Grinvald, & Bonhoeffer, 1997; Yu, Farley, Jin, & Sur, 2005). Given what is known about the structure of these maps and the receptive field properties of the neurons in them, the patterns of activity evoked by specific stimulus combinations will tend to be complex and form a spatially distributed population code. For example, calculations based on known tuning parameters in real orientation maps show that a single small oriented bar stimulus will evoke multiple activity peaks in the cortex whose locations and sizes depend in a complex way on the retinal position and orientation of the stimulus (see Figures 1a to 1g). Likewise, optical recordings from monkey inferotemporal cortex show

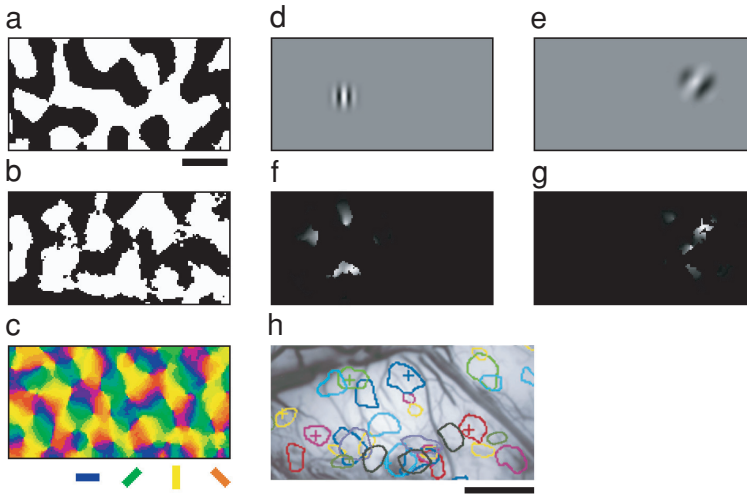

Figure 1: Illustration of population responses of visual cortex to stimuli. (a–c) Examples of cortical map data from area 17 of the cat (data adapted from Hübener et al., 1997). (a) The ocular dominance map (black and white regions respond to the left and right eyes, respectively). (b) The spatial frequency map (black and white regions respond to low and high spatial frequencies, respectively). (c) The color-coded orientation map. All three maps are from the same region of cortex in the same animal. (d, e) Two possible visual stimuli. (d) The stimulus is a small, vertically oriented, high-spatial-frequency Gabor patch seen through the right eye. (e) The stimulus is an obliquely oriented low-spatial-frequency patch seen through the left eye. (f, g) The calculated patterns of activity evoked in the maps by these stimuli. Responses are the product of assumed tuning functions for ocular dominance, spatial frequency, orientation, and spatial position (Swindale, Shoham, Grinvald, Bonhoeffer, & Hübener, 2000). (h) Photograph of area TE of monkey inferotemporal cortex. Regions of cortex shown by optical recording to be activated by a particular visual stimulus (intermediate complexity visual features) are outlined in the same color. Particular stimuli evoke responses in multiple patchy locations. Crosses mark regions where confirmatory electrode recordings were made (reproduced from Wang et al., 1998). Scale bars in *a* and *h* are 1 mm.

that specific stimuli (faces and intermediate complexity features) evoke activity in numerous small (approximately 0.5 mm) spots (Wang, Tanifuji, & Tanaka, 1998; see Figure 1h). Activity patterns in maps that represent stimulus spaces of large numbers of dimensions are predicted to be even more complex (Swindale, 2004). How can such patterns be decoded by higher areas in order to estimate the parameters of the stimulus causing the activity?

While feedforward mechanisms may implement the decoding, in this letter I propose an alternative mechanism in which decoding is done by activating sets of feedback fibers whose axons are tangentially distributed within their target cortical area in such a way that the net spatial activity pattern produced by the fibers models the spatial pattern of feedforward activity in the map. This can equivalently be viewed as a hypothesis-testing procedure or a readout mechanism for a content-addressable memory. Correct decoding, or readout, is signaled when there is sufficient similarity between the feedforward and feedback patterns of activity. This can be measured by summing the net activity within a region of cortex and applying a threshold. The resulting signal could be a single spike or a burst of spikes with a precise temporal relationship to the questioning signal, indicating that the model is a sufficiently good one.

The model follows previous suggestions that the role of feedback pathways is to model or predict patterns of sensory-driven input (Grossberg, 1976; Carpenter & Grossberg, 1987; Mumford, 1992; Rao & Ballard, 1997; Friston, 2005). However, it goes beyond these suggestions by putting them in the context of cortical maps and showing how accurate modeling can be achieved by activating sets of spatially overlapping feedback axons in ways that resemble the combinatorial properties of Venn diagrams (Venn, 1880; Edwards, 2004). The model predicts, or explains, many properties of the cortex, including apical dendrites, the broad tangential distribution of individual feedback axons in V1 (Rockland & Virga, 1989; Rockland & Knutson, 2000; Suzuki, Saleem, & Tanaka, 2000; Shmuel et al., 2005), the finding that attention multiplicatively scales tuning curves (McAdams & Maunsell, 1999; Treue & Martinez-Trujillo, 1999; Martinez-Trujillo & Treue, 2004), and the presence of multiplicative interactions between signals originating in apical and basal dendrites (Larkum, Zhu, & Sakmann, 1999; Larkum, Senn, & Lüscher, 2004; Stuart & Häusser, 2001). Given that cortical map arrangements tend to minimize distances and thus conduction time among functionally related groups of neurons (Durbin & Mitchison, 1990; Chklovskii & Koulakov, 2004), readout should be fast. The model suggests, therefore, that one function of cortical maps may be to minimize readout times.

2 Description of the Model

The receptive fields of individual cortical neurons can be often described as a separable product of functions (often gaussian) of a number of subfeatures, each represented in a spatially organized map. For example a receptive field for orientation, θ, and retinal position, x, y, can be given as the product of three gaussians:

$$f(x, y, \theta) = R_m e^{-\frac{(x-x_0)^2}{2\sigma_x^2}} . e^{-\frac{(y-y_0)^2}{2\sigma_y^2}} . e^{-\frac{(\theta-\theta_0)^2}{2\sigma_\theta^2}}, \tag{2.1}$$

where x_0, y_0, and θ_0 give the receptive field center position and preferred orientation; σ_x, σ_y, and σ_θ give the receptive field tuning widths for position and orientation, respectively; and R_m is the maximum response. This description might not be accurate for some individual neurons (e.g., simple cells), but it is probably a good description of the summed response of a number of cells that are all in a single small column, for example, a minicolumn, 30 to 50 μm in diameter (Mountcastle, 1978). It is only because of this separability that the retinotopic and orientation maps in visual cortex can be considered as separate entities. The assumption also underlies optical recording experiments in which the map for one feature (e.g., preferred orientation) is obtained by averaging over the other features (e.g., spatial frequency, left and right eyes, and retinal position) thought to be represented in the map. These individual feature maps have been termed *protomaps*, and the complete set of protomaps present in a single cortical area has been termed a *polymap* (Swindale, 2000).

Basole, White, and Fitzpatrick (2003) have suggested, on the basis of optical imaging data, that cortical feature maps are not always separable. However, Mante and Carandini (2005) and Baker and Issa (2005) have argued that Basole et al.'s results are compatible with the idea of separable maps if the term *orientation* is understood to mean "the dominant orientation of particular spatiotemporal frequency components of the image." This point will be taken up further in section 6.6.

Separability is important because it makes it possible to model activity patterns in the map produced in response to a particular stimulus as the product of the responses within individual protomaps to each of the particular subfeatures (e.g., orientation, receptive field position) present in the stimulus. Feedback fibers can take advantage of this to model feedforward patterns of activity in a combinatorial way. This can be done by locally multiplying the activities in simultaneously active sets of axons, each of which has a tangential distribution that models the cortical activity pattern produced by a particular value of a subfeature in its corresponding protomap. Figure 2 show a simple hypothetical example of this process in the case of a uniform isotropic retinotopic map and a simulated pattern of orientation domains responding to stimulus at position (x, y, θ) with receptive fields given by equation 2.1. Figure 2a shows the x-retinal protomap response, that is, the net response of the cortex to stimuli at position x, averaged over all other values of y and θ. Figure 2b likewise shows the protomap response to stimuli at position y, averaged across all other values of x and θ, and Figure 2c shows the more familiar spatial pattern of response to an orientation θ, averaged across all values of x and y. Figure 2d shows the predicted response to a single stimulus (x, y, θ), which is simply the product of the activities shown in Figures 2a to 2c.

Three feedback fibers, having tangential weight distributions that match the activity patterns shown in Figures 2a to 2c can model the activity pattern shown in Figure 2d if there is a mechanism that can locally multiply their

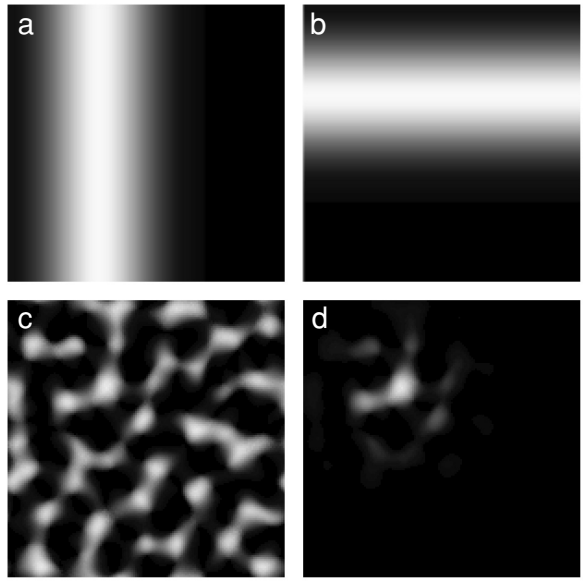

Figure 2: Protomap responses for stimuli at (a) a single retinal x-position, (b) a single retinal y-position, and (c) a single orientation, θ. (d) The product of the activities in a, b, and c. This is the predicted cortical response to a small bar at position (x, y) and orientation θ.

activity. A more plausible mechanism would be to locally add the activities, although this would result in a less accurate model. The consequences of this, and other deviations from the ideal, are not severe and will be examined later. The requirement for local integration could be met by having the fibers converge on apical dendrites, which are able to integrate signals from large numbers of afferent axons at the same tangential cortical location. The resulting patterns of activity originating through feedforward mechanisms would sum, ideally also multiplicatively, with the activity originating in apical dendrites. Net cortical activity will then vary in proportion to the similarity between feedforward and feedback patterns, the total activity being a maximum when the patterns are the same. The net activity could be measured by a type of neuron, here termed a *collector neuron*, whose job is to sum and threshold the activity from all the pyramidal cells within some region of cortex. A correct match between feedback and feedforward patterns would be signaled in an all-or-none way by activity in an appropriately thresholded collector neuron (or neurons) or by determining, via a search procedure, the set of feedback fibers whose activity maximizes the activity of the collector neuron.

Some of the advantages of this scheme as a method of decoding population activity patterns are as follows. Consider a scenario in which the

map shown in Figure 2 is used to encode and decode 10 different retinal x-positions, 10 different y-positions, and 18 different orientations, each 10 degrees apart, yielding a total of 1800 different stimuli in all. It is possible that each pattern could be modeled by a single feedback fiber, with a distribution matching the corresponding pattern of feedforward activity. However, decoding would then require a total of 1800 feedback fibers, while iterative search would require 1800 tests to determine the best match. The alternative combinatorial scheme would require only 38 feedback fibers ($10 + 10 + 18$). Iterative search can also be done much more quickly because the collector response is separable along each of the three dimensions x, y, and θ. Thus, the best fit can be found by finding separately the best values of each of x, y, and θ in turn. This would require at most only 38 tests. While this is an improvement on 1800, the extent to which the brain might use iterative search in visual perception is questionable. First, complex tasks such as visual recognition can often be done very quickly. Second, selecting which of the feedback fibers gives the strongest collector response requires some mechanism for storing intermediate responses and retaining a memory of which fibers produced the strongest ones. However, iterative search is not an essential component of the model. Results will be presented showing how single-shot feedback decoding might be used in two-alternative-forced-choice types of visual discrimination where it is assumed that the brain already has a good internal model of the stimulus.

In the above description, each particular pattern of activity in the feedback fibers explicitly represents a question, or hypothesis, about the nature of the stimulus that is causing a particular pattern of feedforward activity. To simplify the terminology, the feedback fibers will henceforth be referred to as Q-fibers (for "query"). The vector $\mathbf{s} = \{s_1, s_2, \ldots, s_N\}$ will be used to represent the N-dimensional stimulus feature coded by the cortex, where s_1, $s_2 \ldots$ give the values of each of the subfeatures (such as retinal eccentricity and orientation). The spatial pattern of feedforward cortical activity evoked by $\mathbf{s}$ is given by $A(i, j, \mathbf{s})$ where (i, j) index position on a 2D cortical surface. The vector $\mathbf{h} = \{h_1, h_2, \ldots, h_N\}$ will be used to represent a particular hypothesis represented by activity in a subset of the Q-fibers, which is correct when $\mathbf{h} = \mathbf{s}$. Particular values of $h_1, h_2 \cdots$ are presumed to correspond with activity in a particular set of N Q-fibers, each of which contributes an amount $q_n(i, j, h_n)$ to the cortical activity at point (i, j). Figure 3 summarizes the operation of the mechanism.

3 Methods

3.1 Implementation of the Model. Different Q-fiber inputs are assumed to be first integrated locally, by apical dendrites. The integrated signal then modulates the feedforward response, which may be sub- or suprathreshold, and is assumed to originate in inputs to the basal dendrites of the cell. This modulated signal determines the output from the cell. Ideally, all the

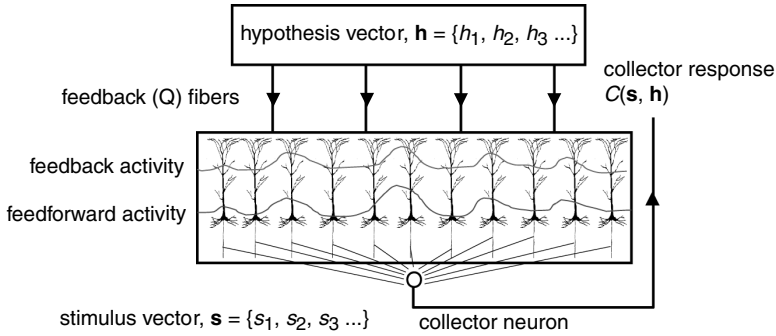

Figure 3: Diagram showing how the decoding mechanism works. A stimulus vector, **s**, produces, via feedforward pathways, a spatial pattern of depolarization (lower line) in pyramidal cells, possibly originating in the basal dendrites. Feedback pathways represent the hypothesis vector, **h**, by producing activity in a set of Q-fibers which terminate on apical dendrites and sum to produce a depolarization (upper line). This then sums with the signal originating in the basal dendrites. The resulting activities of pyramidal cells are summed by a collector neuron, and this signal, which will be largest when $\mathbf{h} = \mathbf{s}$, can be used to decide whether **h** is correct.

Q-fiber inputs are multiplied, together with the feedforward response. Strict multiplication of inputs, however, might be hard to implement neurally. For this reason, several different neural implementations of the model were tested:

1. The additive model: Add all N Q-fiber inputs and the feedforward response linearly, followed by thresholding at a value t. The output O at point (i, j) in the map is given by

$$O(i, j) = thr \left\{ A(i, j, \mathbf{s}) + \sum_{n=1}^{N} q_n(i, j, h_n) \right\}. \tag{3.1}$$

2. The add-multiply model: Add all the Q-fiber inputs, but assume that the sum multiplicatively scales the feedforward response, followed by thresholding:

$$O(i, j) = thr \left\{ A(i, j, \mathbf{s}) \sum_{n=1}^{N} q_n(i, j, h_n) \right\}. \tag{3.2}$$

3. The multiplicative model: Multiply all the inputs:

$$O(i, j) = A(i, j, \mathbf{s}) \prod_{n=1}^{N} q_n(i, j, h_n). \tag{3.3}$$

Further possibilities exist, including passing the summed Q-fiber inputs and then the summed Q and A inputs through sigmoidal nonlinearities before calculating O. Although some of these models have been implemented, the description of results will be restricted to these three models for simplicity.

The collector response $C(\mathbf{s}, \mathbf{h})$ is given by

$$C(\mathbf{s}, \mathbf{h}) = \sum_{i,j} O(i, j). \tag{3.4}$$

As discussed above, $\mathbf{s}$ can be estimated by maximizing C separately for each of the dimensions of $\mathbf{h}$, taken in turn. However, it seems unlikely that the brain normally engages in a search of this nature. More realistic, perhaps, is a single-shot scenario, where a hypothesis is generated and a threshold is applied to the collector signal to generate a yes-no response. For a given hypothesis and a suitably chosen threshold, estimates can then be made of the change in stimulus value, which will lead to correct discrimination on a criterion percentage of trials, as in a two-alternative forced-choice psychophysical experiment.

Calculations were done to study the accuracy with which $\mathbf{s}$ can be estimated. Factors that may affect this include (1) the particular type of model (additive, multiplicative or add-multiply); (2) noise in the feedforward response; (3) the particular type of map, for example, with good or poor coverage uniformity; (4) whether the stimulus is presented in the same or slightly different retinal locations in successive tests; (5) the dimensionality of the map; and (6) the tuning widths of the receptive fields.

3.2 Map Generation. Estimates of decoding accuracy were made using artificially generated maps of feature spaces that included a two-dimensional "retina" and N additional angular dimensions. Because there is no evidence for independent cortical maps of more than one angular variable, most of the calculations were restricted to maps for which $N = 1$. However, some were done for $N > 1$ in order to show the generality of the model. The model was additionally tested on fully or incompletely populated binary feature maps of six or eight dimensions. Maps were generated using the Kohonen self-organizing feature map algorithm (Kohonen, 1982). Implementation details were similar to those described previously (Swindale, 2000, 2004). Maps were 150×150 pixels in size, the retina was 12×12 units in size, periodic boundary conditions were implemented, and parameters were adjusted to give protomap periodicities of around 28 pixels, corresponding roughly to 1 mm in the real cortex. Maps were generated with, and without, simulated annealing. Details of the annealing schedule are given in Swindale (2000, 2004). Annealing—gradual reduction of the width of the cortical neighbourhood function in the Kohonen algorithm during development—tends to produce a more uniform distribution of

stimulus features (i.e., more uniform coverage) across the map. In order to examine the effect of coverage uniformity on decoding accuracy, the cortical neighbourhood function was reduced in size by different amounts during development to produce a series of maps varying in coverage uniformity.

3.3 Estimates of Decoding Accuracy. The cortical activity pattern $A(i, j)$ produced in response to a stimulus, $\mathbf{s} = \{x, y, \theta_1, \theta_2, \ldots, \theta_N\}$ was given by

$$A(i, j) = f(\mathbf{s} - \mathbf{w}_{i,j}), \tag{3.5}$$

where $\mathbf{w}_{i,j}$ is the feature mapped to point (i, j) in the cortex. The receptive field shape, f, was assumed to be constant over the cortex and was given by a product of gaussians,

$$f(x, y, \theta_1, \theta_2, \ldots, \theta_N) = e^{-(x^2+y^2)/2\sigma_r^2} \prod_{n=1}^{N} e^{-\theta_n^2/2\sigma_\theta^2}, \tag{3.6}$$

where σ_r and σ_θ give the retinal and angular receptive field sizes, respectively.

Multiplicative noise in the feedforward response was modeled by setting

$$A'(i, j) = A(i, j)[1 + \xi(i, j)], \tag{3.7}$$

where ξ is a random number drawn from a gaussian distribution with a standard deviation $= \sigma_A$. Negative values of A' were permitted on the grounds that A might represent not the thresholded spiking response of the cell but the state of membrane depolarization of the soma induced by feedforward inputs, with which signals originating in the apical dendrite interact.

Q-fiber density (or net postsynaptic effect) at a given point in the cortex was assumed to be a gaussian function of the difference between the hypothesized subfeature value and the value represented in the map at that point. Thus, if the feature value represented at point (i, j) in the cortex is $\mathbf{w}_{i,j} = \{x_c, y_c, \theta_{1c}, \theta_{2c}, \ldots, \theta_{Nc}\}$ and the hypothesis vector $\mathbf{h} = \{x_h, y_h, \theta_{1h}, \theta_{2h}, \ldots, \theta_{Nh}\}$, retinal Q-fiber density, q_r, was assumed to be given by

$$q_r(i, j, x_h, y_h) = e^{-\frac{(x_h - x_{i,j})^2 + (y_h - y_{i,j})^2}{2\sigma_r^2}}. \tag{3.8}$$

Note that the density distribution in the cortex will be gaussian only if retinotopy is uniform, which is generally not the case for the maps studied

here. Also, unlike the example given above (see Figure 2), in this formulation a single Q-fiber codes for both horizontal and vertical components of stimulus position. This was done partly for simplicity and partly because there is little evidence to suggest that top-down (e.g., attentional processes) can independently target x- and y- components of stimulus position.

The Q-fiber density for each of the $n = 1$ to N orientation dimensions was given by

$$q_n^{\theta}(i, j) = e^{-\frac{(\theta_{n,h} - \theta_{n,i,j})^2}{2\sigma_r^{h2}}}, \tag{3.9}$$

where $\theta_{n,i,j}$ is the value of the nth orientation parameter at point (i, j) in the cortex and $\theta_{n,h}$ is the value of the nth hypothesis orientation.

The parameters controlling the distribution widths of the Q-fibers, σ_r^h and σ_θ^h, were generally, although not always, assumed to have the same widths as the corresponding feedforward tuning functions. When that is the case, the product of the set of $N + 1$ Q-fibers corresponding to a particular hypothesis, **h,** would produce an activity pattern exactly equal to the pattern of activity produced by the feedforward response to a stimulus $\mathbf{s} = \mathbf{h}$. This happens only in the multiplicative version of the model (see equation 3.3). It is assumed that a learning mechanism allows the distributions of individual Q-fibers to correctly match the different components of the feedforward response. How this might work will not be considered further here. It is also assumed that although many Q-fibers might innervate a cortical region, only those sets that correspond to a specific hypothesis are simultaneously active during decoding, that is, that the hypothesis is well formulated.

4 Results

4.1 Orientation Decoding. Figure 4 gives an example of how decoding works for an orientation map (see Figure 4a). Figure 4b shows an evoked pattern of feedforward activity, assuming physiologically realistic receptive field size parameters. The stimulus is positioned in the middle of the model retina, which maps approximately to the center of the cortex. Figure 4c shows the corresponding retinal Q-fiber distribution calculated according to equation 3.8. The irregular distribution is caused by local irregularities in the retinotopic map. Figure 4d shows the Q-fiber distribution corresponding to the stimulus orientation. Figure 4e shows the activity pattern in the map as modified by the Q-fiber activity. Because the add-multiply model was used, the pattern is obtained by adding the two Q-fiber activities (see Figures 4c and d), multiplying by the feedforward pattern (see Figure 4b), and thresholding to obtain the pattern shown in Figure 4e. The collector response, C, is the sum of the activity values in Figure 4e taken across the whole map.

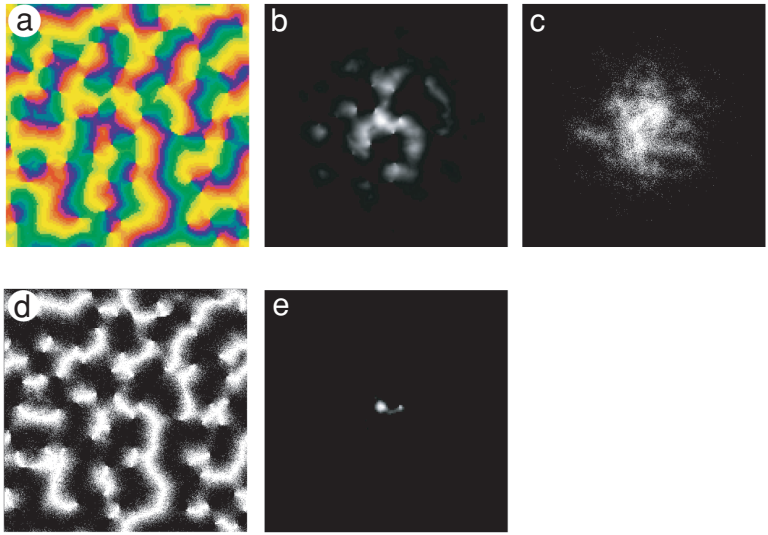

Figure 4: Illustration of the decoding mechanism for orientation maps. An add-multiply model with a threshold $t = 1.0$ was assumed. (a) A color-coded orientation preference map. (b) The calculated feedforward activity pattern evoked by a stimulus $\mathbf{s} = \{6.0, 6.0, 0.0\}$, assuming an orientation tuning width of 25 degrees and a retinal receptive field size of 1.12 units. (c, d) The calculated Q-fiber distributions for the matched hypothesis $\mathbf{h} = \{6.0, 6.0, 0.0\}$ for the retinal and orientation parameters, respectively. (e) The pattern of activity, $O(i, j)$ (see equation 3.2), produced by summing the Q-fiber activities (c, d), multiplying by the feedforward inputs (b), and thresholding. The collector response, C, on which decoding is based, is the sum of all the activity values in e.

Figure 5 shows how the collector response varies as a function of stimulus orientation, θ_s, and the hypothesis orientation, θ_h, for the same map and model as shown in Figure 4, for a fixed retinal stimulus position and no response noise. Note that the value of θ_h for which C is a maximum is usually a few degrees away from θ_s. This is one possible estimate of the accuracy of decoding, although iterative search is required. In addition, the values of C for $\theta_s = \theta_h$ vary. This is a consequence of uneven coverage—the fact that different stimuli evoke slightly different net responses. As discussed below, this variability is likely to influence the accuracy of readout.

Figure 6 shows how the distribution of values of C changes as a function of the difference between θ_h and θ_s. This suggests a basis for an alternative, one-shot estimation of stimulus parameters. One can ask what the smallest difference is between two stimuli that changes the value of C by an amount sufficient to support a yes-no decision about the stimulus, as in a psychophysics experiment. The size of this difference was estimated

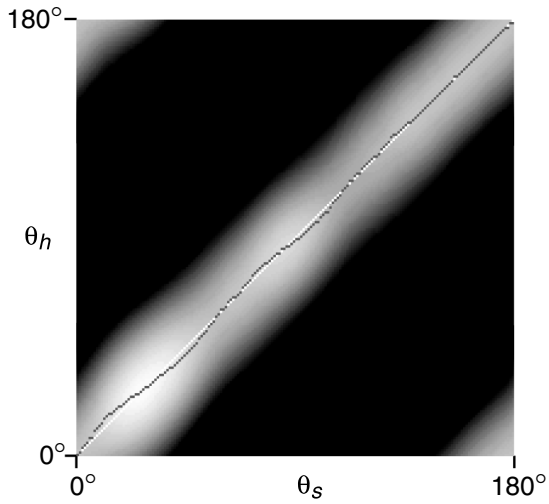

Figure 5: The magnitude of the collector response, C, as a function of θ_s and θ_h (cf. Figure 3e for the discrete case) for a one-dimensional orientation map. The white line shows the locus of points for which $\theta_s = \theta_h$, and the black points near it indicate, for each value of θ_s, the value of θ_h for which C is a maximum. The add-multiply model was used to calculate the values of C, with a threshold $t = 1.0$. All stimuli were presented at a single retinal location, $x_s = y_s = 6.0$, and it was assumed that the decoding mechanism knew the retinal location, so that $x_h = y_h = 6.0$.

by measuring the mean, $C_{mean}(0)$, and standard deviation $C_{sd}(0)$, for a set of randomly chosen stimuli for which $\mathbf{s} = \mathbf{h}$. The smallest orientation difference, $\Delta\theta_t$, for which $C_{mean}(0) - C_{sd}(0) > C_{mean}(\Delta\theta) + C_{sd}(\Delta\theta)$ was then estimated by calculating $C_{mean}(\Delta\theta)$ and $C_{sd}(\Delta\theta)$ at 1 degree intervals. This criterion gives a threshold, $C_{mean}(\Delta\theta_t)$, which, if applied to the value of C, will yield 84% correct identification, similar to the rate used in many psychophysical experiments.

4.2 Effects of Coverage and Response Noise. Orientation thresholds, measured as just described, are limited by the fact that stimuli are presented at randomly chosen retinal locations within the map and, because of nonuniform coverage, evoke slightly different net responses. A second factor likely to influence thresholds is the intrinsic variability in the responses of individual neurons to stimuli. This noise is typically multiplicative rather than additive in character. The relative influence of coverage and response noise was studied by measuring thresholds as a function of (1) coverage uniformity in the absence of response noise and (2) multiplicative noise levels in situations where stimuli were presented at random retinal locations

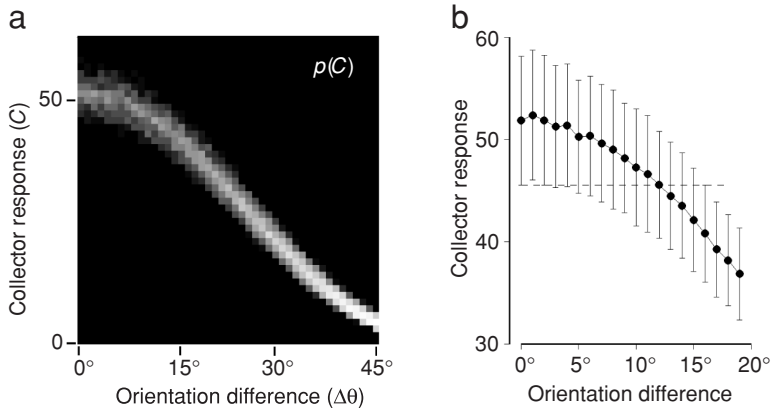

Figure 6: (a) The probability distribution of values of the collector response C, as a function of the difference between θ_s and θ_h in one-degree steps for stimuli presented at randomly chosen retinal locations and a fixed stimulus orientation $\theta_s = 90$ degrees. Values were obtained from an $N = 1$-dimensional orientation map with an add-multiply model with a threshold $t = 1.0$. It was assumed that the decoding mechanism knew the retinal location, so that $x_s = x_h$ and $y_s = y_h$. Probability values were based on 1000 presentations at each orientation difference value ($\Delta\theta$) with no response noise. (b) Illustration of the method used to determine a one-shot threshold: the graph shows the mean and standard deviations of the values of C obtained as in a. Threshold is defined as the smallest orientation difference for which the error bars (=1 SD) fail to overlap. The dashed line shows the threshold value of C, which, if used as a basis for a yes-no decision, would produce an 84% correct detection rate. In this case the threshold is about 16 degrees.

or at a fixed location. In the latter case, response noise should be the only factor limiting thresholds.

A measure of coverage uniformity (Swindale, 1991) is $c' = \frac{c_{sd}(s)}{c_{mean}(s)}$ where the mean and SD of C are taken across some representative set of stimuli. In the absence of other sources of noise, coverage-related variability in the values of $C(s)$ would be the sole determinant of threshold. This was demonstrated by generating a series of orientation maps that differed in their coverage (c') values but were similar in other respects. As expected, there was a strong correlation between c' and thresholds for orientation (see Figure 7a) and retinal position (see Figure 7b). Figure 7c shows how orientation thresholds vary as a function of multiplicative response noise (see equation 3.7) for stimuli presented at random (filled circles) or fixed (open circles) positions. For random positions, thresholds should be limited by a combination of coverage noise and response noise. However, since these are likely to add quadratically, low levels of response noise should

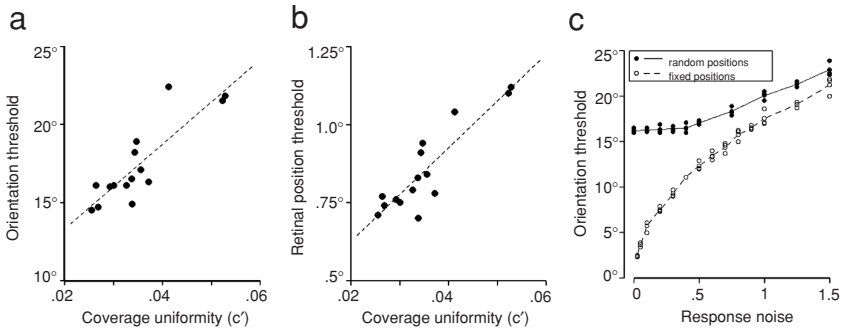

Figure 7: The relationship between one-shot thresholds for (a) orientation and (b) position, as a function of coverage uniformity. Each point is a measurement from a single map made using the multiply model. Similar results were obtained with the other two models. (c) The effect of multiplicative response noise on orientation discrimination thresholds for stimuli presented at randomly chosen retinal locations (filled circles) or at a fixed location (open circles). The orientation of each stimulus was fixed at $\theta_s = 90$ degrees, and the multiply model was used. The x-axis shows the value of σ_A (see equation 3.7); symbols show the thresholds calculated for three different maps, and lines connect the means of each set.

have little impact on thresholds until their magnitude is comparable to that of coverage noise. This is evident in Figure 7c, where levels of response noise below about 0.5 have little impact on thresholds. For fixed positions, where coverage noise is absent, thresholds approach zero as noise levels approach zero.

Because of microsaccadic eye movements, most psychophysical measures of orientation discrimination are likely to involve some variability in the retinal location of a briefly presented stimulus. For that reason, it was decided to calculate thresholds on the basis of randomly chosen stimulus locations. Because coverage noise appears to be the major factor limiting thresholds in this case, response noise was assumed for simplicity to be zero, since even quite high levels have little impact on thresholds (see Figure 7c).

4.3 Other Factors Affecting Thresholds. Figure 8a shows how the one-shot threshold, $\Delta\theta_t$, varies as a function of the number of orientation dimensions for the add-multiply, multiply, and add-threshold models. In all of the tests, the retinal position of the stimulus was varied randomly and $\mathbf{h} = \mathbf{s}$ with the exception of a change in only a single orientation value, $\theta_{1,h}$. Thresholds rise with the number of orientation dimensions, and, as expected, the multiply model generally performs best. While thresholds are poor for $N > 2$ orientations, the one-shot method used here, although fast, is probably the least good decoding method that could be used; other

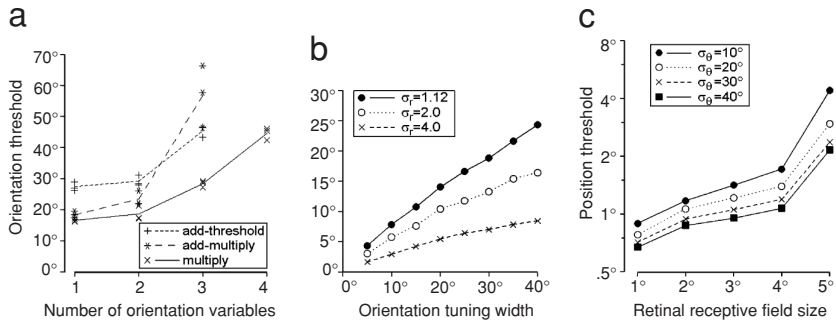

Figure 8: (a) Estimates of orientation discrimination threshold as a function of the number of orientation dimensions (N) represented in the map for the three different models. For the add-threshold model, $t = N + 0.5$; for the add-multiply model, $t = 0$. Matched tuning widths for the feedforward and Q-fiber pathways were assumed; for orientation, $\sigma_\theta = 25$ degrees, and for retinal receptive fields, $\sigma_r = 1.12$. Symbols show the thresholds calculated for each of three different maps, and lines connect the means of each set. (b) Orientation thresholds as a function of orientation tuning width, for three different retinal receptive field sizes. (c) Retinal position thresholds as a function of retinal receptive field size for four different orientation tuning widths and a fixed stimulus orientation. Thresholds shown in b and c are from maps with $N = 1$ using the multiply model; matched tuning widths for feedback and feedforward pathways were also assumed. Points show the mean of the thresholds in three maps. Similar, though slightly less good, values were obtained using the add-multiply or add-threshold models.

readout methods could probably be devised that would give better results. Figure 8b shows how orientation tuning width affects orientation discrimination for a variety of retinal receptive field sizes. Better orientation discrimination is obtained with narrower orientation tuning and larger retinal receptive fields. Figure 8c shows that for retinal position, better discrimination is obtained with narrower retinal tuning and broader orientation tuning.

4.4 Effects of Differences in Feedforward and Feedback Tuning Widths. So far, it has been assumed that the tuning widths that characterize the feedforward receptive field and the spatial distribution of the feedback fibers are the same. It might be expected that unequal tuning widths per se would lead to less accurate decoding, since the model of feedforward activity will be less good. However since smaller tuning widths generally lead to more accurate decoding, accuracy might be improved if the width of feedback, or feedforward, tuning is reduced while keeping the other constant, but only up to a point. In addition, because the multiplicative model

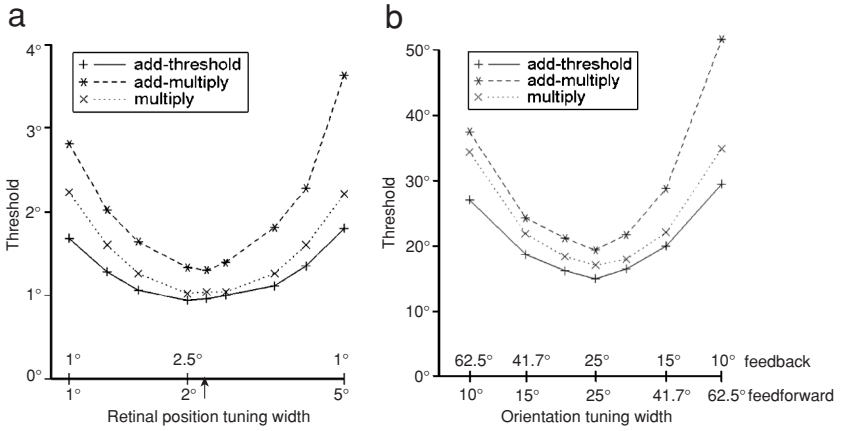

Figure 9: The effect of changing the widths of feedforward and feedback tuning functions on discrimination thresholds. Widths were varied reciprocally, so that the product of the two widths was always the same. (a) Results for retinal position tuning width: arrow indicates equality of the two parameters ($=\sqrt{5}$); (b) results for orientation tuning. A threshold of $t = 2.5$ was used for the add-threshold model in a and b.

is symmetric with respect to feedforward and feedback tuning widths, it will not matter which of the two is changed. For the other models, this would not be the case.

These expectations were tested by reciprocally varying the tuning widths of feedforward and feedback retinal receptive fields (Figure 9a) or of orientation tuning widths (Figure 9b) so that the product of field sizes remained constant. As expected, one-shot decoding thresholds were lowest when fields had the same size. Calculations confirmed that thresholds were unchanged by exchange of unequal feedforward and feedback tuning widths for the multiplicative model. For the other two models, there was a slight advantage, under exchange, in having feedforward tuning be the narrower of the two widths. If the tuning width of one parameter was kept constant and the other was narrowed, thresholds decreased and then started to increase, the increase reflecting the effect of the increasing inequality in the widths of the tuning functions.

4.5 Effects of Apical Dendritic Spread. Accurate reconstruction of the feedforward activity pattern requires local integration of Q-fiber inputs at each point in the map. This integration can be performed by apical dendrites that sum many inputs within a small ($<50\ \mu$m) tangential region of cortex over most of their length. To study how important a narrow integration region is to the performance of the model, the effect of lateral

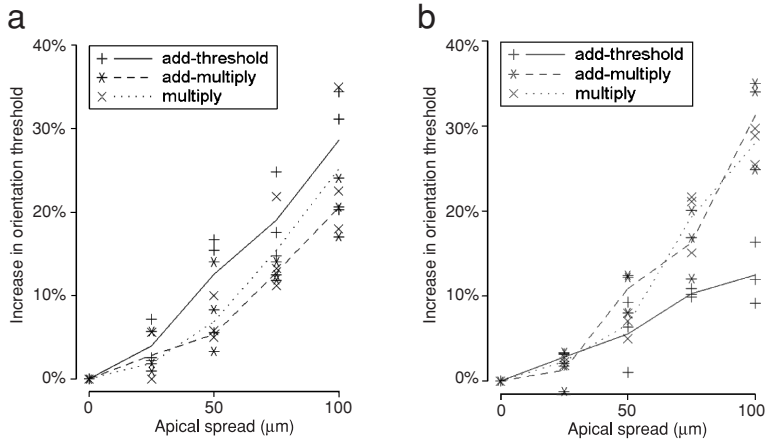

Figure 10: The effect of lateral spread of the integration regions of apical dendrites on orientation discrimination thresholds for a maps of (a) $N = 1$ or (b) $N = 2$ orientation variables. Symbols show the percentage change in thresholds compared to the threshold for no spread ($\sigma_d = 0$) calculated for each of three different maps using each of the three different models. Lines connect the means of each set of three measurements for each type of model.

spread of apical dendrites on thresholds was studied. This was done by convolving each of the Q-fiber distributions with a gaussian of width σ_d, thus simulating an equivalent gaussian spread of integration by the apical dendrites. Measurements of one-shot orientation thresholds were made as described above for different values of σ_d. Figure 10 shows the percentage change in threshold for orientation maps of $N = 1$ (see Figure 10a) or 2 (see Figure 10b) for the three models, as a function of apical spread. Thresholds rise by about 1% to 3% for every additional 10 μm of spread in each of the three models.

5 Binary Decoding

Calculations, not presented here, were done to show that the model can be extended to binary feature maps representing discrete parameter values (Swindale, 2000). In this setting, feedback fibers have distributions that are spatially all-or-none, rather than graded in strength, and the distribution of particular Q-fibers matches the spatial representation of the corresponding binary subfeature value (0 or 1) in the map. Detection of the presence of a particular binary feature is accomplished by activating the set of N Q-fibers corresponding to the N bit values of the feature being detected. Regions where all N fibers overlap correspond to regions of cortex representing

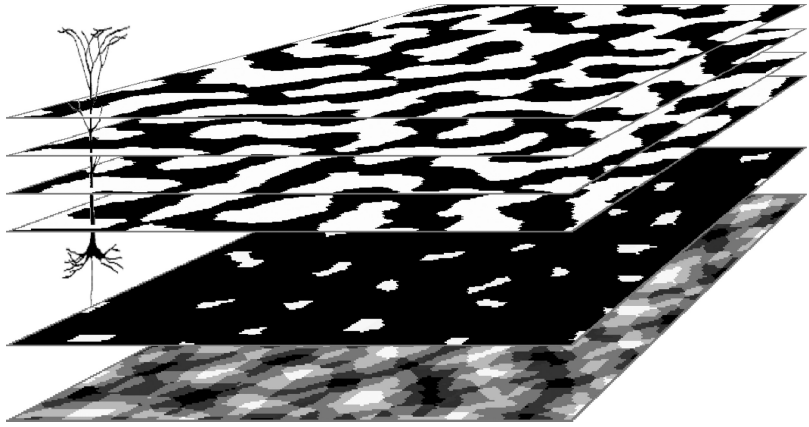

Figure 11: Illustration of how decoding works for binary features. The top four layers show, in white, the tangential distributions of four feedback axons (Q-fibers), which overlap like the subsets in a Venn diagram. Regions of overlap (=logical conjuctions) can be detected by summing activity in the fibers with the apical dendrites of pyramidal neurons. The bottom two layers show the resulting activity in a sheet of pyramidal cells if the Q-fiber activities are summed linearly (lowest layer) or multiplicatively (second lowest layer).

the feature, and the conjunction can be detected by summing with apical dendrites (see Figure 11), and then summing the outputs of pyramidal cells with collector neurons, as for the continuous model. As with the continuous case, a map representing a binary space of N dimensions could represent 2^N features, but only a total of $2N$ Q-fibers would be required to do the decoding. Additional fibers can be used to decode retinal position, in the same way as for orientation maps. The model has been tested with binary feature maps of varying numbers of dimensions and with an appropriately chosen threshold on the collector response is capable of achieving 100% recognition accuracy for $N \leqslant 8$ dimensions.

It is interesting to consider the case where not all 2^N binary features are represented in the map. The absence of a binary feature from a region of the map corresponds to the geometrical property that the corresponding set of Q-fibers fails to have an area of common overlap anywhere in the region. In this setting, the model functions as a content-addressable memory. Two functions of the mechanism are possible: (1) to ask whether a particular binary stimulus is physically present and (2) to ask if a particular feature is represented in the map, even though not physically present. The latter possibility may give clues to how visual cortex might function when the eyes are closed and feedforward inputs are absent (Zangaladze, Epstein, Grafton, & Sathian, 1999; Kosslyn & Thompson, 2003).

6 Discussion

There are three main components to the ideas presented in this letter. The first is that feedback pathways attempt, in a combinatorial way, to model, or reconstruct or predict, spatial patterns of activity in the areas to which they project; the second is that apical dendrites play a critical role in the reconstruction; and the third is the idea that the goodness of the match is measured by summing across a population of pyramidal cells and thresholding to create a signal that may be used as a basis for subsequent action. These three components will be considered further in the following sections.

6.1 Reconstruction by Feedback Pathways. Many authors have proposed that cortical feedback is used to reconstruct, or predict, patterns of feedforward activation (Grossberg, 1976; Carpenter & Grossberg, 1987; Mumford, 1992; Hinton, Dayan, Frey, & Neal, 1995; Dayan, Hinton, Neal, & Zemel, 1995; Olshausen & Field, 1997; Rao & Ballard, 1997; Lee & Mumford, 2003; Friston, 2005). In the Bayesian framework within which some of these models are constructed, feedback activity is the brain's estimate of the prior probability of a particular neural signal, and the feedforward activity is the posterior probability. The product of the two gives the actual probability and requires neural multiplication of much the same kind that is postulated here. Models based on predictive coding (Mumford, 1992; Rao & Ballard, 1997) suggest that prediction is used to minimize brain activity by subtracting the predictive signal from the one produced by direct sensory feedforward activation. Although the idea that the brain tries to minimize its own activity is attractive, the hypothesis does not seem to be well supported by neurophysiological evidence, which suggests that feedback is usually facilitatory. For example, attending to sensory stimuli almost always enhances responses (McAdams & Maunsell, 1999; Treue & Martinez-Trujillo, 1999; Chawla, Rees, & Friston, 1999; Reynolds, Pasternak, & Desimone, 2000; Martinez-Trujillo & Treue, 2004), often by multiplicative scaling of tuning curves. The present model could, if required, be recast in a form in which the feedback signal was reconstructed in a separate population of pyramidal cells not directly driven by feedforward inputs, which then inhibited, via interneurons, the feedforward signal. However, this would remove much of the simplicity from the model and would not require individual cells to have both tangentially and narrowly spreading dendrites.

While it seems natural to equate Q-fiber activity with top-down, attention-based processing, it is possible that feedback is important in processing unattended as well as attended stimuli. For example, experiments on anesthetized animals, in which attentional mechanisms are presumably inoperative, show that inactivating sources of feedback from higher visual areas to lower ones leads to a depression of stimulus-evoked activity in the lower areas (Hupé et al., 1998; Wang, Waleszczyk, Burke, & Dreher, 2000;

Galuske, Schmidt, Goebel, Lomber, & Payne, 2002; Huang, Chen, & Shou, 2004). Thus, feedback mechanisms, and by implication the mechanism proposed here, may be operative during all forms of visual processing.

6.2 Role of Apical Dendrites. Although computational models of feedback behavior in the brain have taken many forms, none to the author's knowledge has addressed the problem in the context of the multidimensional nature of receptive fields and the fact that patterns of activity in the brain in response to stimuli are often spatially structured because of the presence of topographic maps. It is because receptive fields can be expressed as separable products of functions of individual stimulus parameters (such as orientation and position) that it is possible to model the spatial patterns of activity produced by a multidimensional stimulus feature as the product of patterns of activity due to each of the subfeatures. This allows the very large number of possible patterns of feedforward activity to be modeled in a combinatorial fashion by a relatively small number of feedback fibers. Accurate reconstruction of the pattern, however, requires a spatially local integration of the inputs, which, it is hypothesized here, is the function of apical dendrites. Integration of the activities of individual feedback (Q-fiber) inputs should ideally be multiplicative, and this signal should also combine in a multiplicative way with the feedforward signal originating in the basal dendrites. Although the calculations done here show that accurate decoding can be obtained with linear summation of signals followed by thresholding, in almost all cases multiplicative interactions produced more accurate decoding (see Figure 8a). Thus, neural mechanisms implementing multiplicative interactions between apical and basal inputs should have a selective functional advantage and would be expected to have evolved.

Neurophysiological evidence does indeed show that inputs to the apical dendrites of pyramidal neurons multiplicatively scale responses originating in the basal dendrites, providing they are timed to coincide with a backpropagating action potential (Larkum et al., 1999; Larkum, Senn, & Lüscher, 2004; Stuart & Häusser, 2001) within approximately ± 5 msec. The same authors show that this interaction may also trigger bursts in pyramidal cell firing. These bursts could provide a signal that might be selectively detected by collector neurons. Burst mode firing might thus be associated with a recall or decoding state of the cortex separate from one that reflected mainly feedforward processing.

6.3 Collector Neurons. Neurons that provide a readout, or match signal, by summing local activity are an essential, but speculative, part of the model. It would not be surprising, however, if such cells went unnoticed by physiologists. First, only a small percentage of cortical neurons would need to be collector neurons, so they might not be encountered very often. Second, they might be silent much of the time (if readout was not occurring) or, if active, would have large and unselective receptive fields.

These behaviors would probably be dismissed as of little interest in any sample of recorded cells. However, there is some indirect evidence for the existence of neurons with the properties of collector neurons. It is known that a component of the optical response to visual stimuli in primary visual cortex is stimulus driven but is unspecific to the particular stimulus used (Bonhoeffer & Grinvald, 1996). It is possible that this response reflects the activities of collector neurons. In a population of 40 simultaneously recorded retinal neurons responding to natural scene stimuli, Schneidmann, Berry, Segev, & Bialek (2006) reported the presence of "check" cells, whose probability of firing at any instant was linearly proportional to the total number of spikes present in the rest of the network. This is exactly the property required of collector neurons. While these cells are unlikely to be performing the function postulated here (since the retina receives no feedback fibers), the techniques used by Schneidmann et al. (2006) might be used to identify collector neurons in simultaneous recordings from cortical areas.

6.4 The Readout Mechanism. The one-shot method used here is simple and could be elaborated on. It is meant to be a model for decoding in two-alternative forced-choice experiments where the subject has to decide into which of two categories a stimulus falls within a very brief period following presentation of the stimulus. A threshold is applied to the collector response; the response is yes if the threshold is exceeded and no if it is not. Orientation discrimination thresholds for small bars calculated with this method (see Figure 8) are in the range of 5 to 20 degrees (depending on the precise widths of the receptive fields involved). These values may seem poor, but they may not need to be any better given that the information present in natural images can be well represented by channels with an orientation resolution of about 10 to 30 degrees (van der Schaaf & van Hateren, 1996). These thresholds are also comparable to those obtained by Watt (1987) in humans for a very short, briefly flashed bar, which is the appropriate comparison stimulus. Thresholds were about 30 degrees for a bar 3 arc minutes in length (comparable in size to a foveal receptive field) presented for 35 ms and decreased to about 8 degrees for a longer presentation time of 100 ms.

A disadvantage of the fixed threshold mechanism is that it will fail if stimulus parameters other than orientation (or the one being tested) are varied between presentations. Thus, if the contrast or the length of the bar is changed at random, the magnitude of the collector response will change in a way that cannot be disambiguated from a change caused by the stimulus parameter being discriminated. Most psychophysical estimates of threshold are obtained under carefully controlled conditions in which only the tested dimension is varied. Randomly varying the spatial position of stimulus presentation (Cohn & Lasley, 1974; Lindblom & Westheimer, 1992) or even making the time of presentation uncertain (Westheimer & Ley, 1996) can elevate orientation discrimination thresholds significantly. Thus,

uncertainty about stimulus parameters other than the one being tested can elevate thresholds. However, a fixed threshold mechanism would probably degrade even more severely than this. Other readout methods might give better performance and be less susceptible to stimulus uncertainty. In the limit, an iterative search leading to the identification of the hypothesis maximizing the collector response (gradient ascent on the collector response in **h**-space) would lead to stimulus identification that was not affected by uncertainty (including contrast) in any of the parameters. Thresholds obtained this way (data not shown) are substantially smaller than those obtained with the one-shot method. However, iterative search might be hard to implement neurally and would also be slow. A compromise might involve testing a few strategically chosen hypotheses, for example, one at the midpoint of the expected stimulus range and two on either side. This mechanism would require temporary storage of collector responses and the ability to make a yes-no decision based on the stored magnitudes. For these reasons, it might be of interest to examine the interactions between stimulus uncertainty, presentation time, and reaction time on thresholds in human subjects.

Hypotheses that are wrong in one or more feature dimensions would lead to small collector responses, and this could be used to explain why subjects can be unaware of highly salient stimuli when attentional (that is, hypothesis-generating) resources are directed elsewhere. On the other hand, it is obvious that when attentional resources are not limited, one can see things that are not expected. In this case, it seems reasonable to presume that the initial wave of feedforward activity following presentation of a stimulus is able, most of the time, to drive the generation of predictions that are correct or close to correct.

6.5 The Q-fibers. Feedback fibers constitute a high proportion of all connections in the brain (Salin & Bullier, 1995). Conduction velocities are similar to those of feedforward axons (Girard, Hupé, & Bullier, 2001) and effects can be rapid (Pascual-Leone & Walsh, 2001; Hupé et al., 2001). Since large numbers and high speed must be achieved at the expense of brain volume, which is at a premium, the function of the brain must be critically dependent on fast feedback as well as fast feedforward connectivity. This argues strongly against the widespread notion that feedback connections subserve a slow modulatory role or that they are mainly involved in attentional processes that operate on a timescale of hundreds of milliseconds. It suggests that the normal communication of information between brain areas, which operates on a scale of milliseconds or tens of milliseconds, requires the continuous fast integration of both feedforward and feedback activity.

The particular mechanism suggested here implies a significant distinction between the types of information transmitted backward and forward. Feedforward signals constitute a population code, in which the significance

of a firing rate in a particular neuron is ambiguous and can be interpreted only within the context of the overall pattern of firing within a large group of cells (Pouget, Dayan, & Zemel, 2000). Thus, the information represented is implicit rather than explicit. In contrast, the feedback information represented in the Q-fibers constitutes an explicit code. That is, synchronous spikes in a particular set of Q-fibers form an explicit representation of a stimulus (or more accurately constitute a mental event modeling the stimulus), and the pattern corresponds uniquely to that event. The firing of an individual Q-fiber has the property that it should be invariant over a large class of stimuli (or events) but strongly dependent on the presence of a single subfeature (or subevent). The nature of the subfeature would, of course, depend on the type and level of the representation.

There is little physiological evidence that directly supports the existence of Q-fibers. This may not be surprising as physiologists generally focus their attention on neurons that are selective rather than unselective in their responses to stimulus features. Anatomical evidence, however, supports the prediction that feedback fibers should have a wide tangential distribution matching the layout of functional response properties (or subfeatures) represented in the map to which the fiber projects. For example, a Q-fiber representing a given orientation angle would have a widespread patchy distribution of terminals whose density (or net weight) matched the magnitude of the map response to a stimulus with that angle, averaged over all other stimuli (such as spatial frequency, ocular dominance and retinal position). A Q-fiber representing spatial position would have a somewhat narrower and more continuous tangential distribution of inputs (e.g., as shown in Figure 4c). In reconciling these expectations with the anatomical data on feedback projections it should be borne in mind that (1) Q-fibers might be only a subset of feedback fibers; (2) although it seems likely, the model does not require that Q-fibers be cortical in origin, for example, they might originate in the thalamus or other subcortical nuclei; and (3) what is described here as the distribution of inputs from a single fiber might in reality be the sum of a number of fibers that normally act in synchrony. It might then be quite hard to relate the distribution of a single axon to a particular functional property of the map within which it arborized.

Bearing these qualifications in mind, many of the properties of cortico-cortical feedback axons in the cortex match the predictions of the model. In general, feedback projections of individual axons spread further tangentially than do feedforward connections, they often have a patchy distribution, and they are usually present in the upper layers, where they can make contact with apical dendrites (Rockland & Virga, 1989; Rockland & Knutson, 2000; Suzuki et al., 2000). Feedback projections from V2 to V1 have a patchy distribution of terminals whose local density depends on the difference between the orientation preference of the region of origin in V2 and the orientation preference of the region of V1 concerned (Shmuel et al., 2005), exactly as postulated here for Q-fibers that represent orientation

(see equation 3.9). It is not known whether other feedback afferents to V1 (e.g., from areas MT and V4) have distributions that match other mapped functional properties such as color or retinal position.

Two further related questions concerning Q-fibers are not answered here. First, nothing is said about how valid hypotheses—patterns of firing that potentially can model an input—might be generated. For example, many patterns of firing would not correspond to any possible feedforward input. If, as seems likely, Q-fiber inputs to a given cortical area come from a variety of other areas, a mechanism is needed to coordinate the firing of the relevant cells in the different areas. Subcortical mechanisms might be required to do this. Extending the model hierarchically might lead to a better understanding of how this could be done. Second, it needs to be demonstrated that the Q-fiber distributions can be learned by a plausible mechanism. To do this, the fibers need to factorize the input distributions (which are represented as the product of activity patterns). Mechanisms for doing this can be envisaged, but a detailed implementation is beyond the scope of this letter.

6.6 Map Separability. A fundamental assumption of the model is that cortical receptive fields can be described as separable products of functions along several dimensions, each of which is represented in a spatially ordered map. This makes it possible to determine, for each point in the map, best values of each parameter, for example, orientation, spatial frequency preference, ocular dominance and retinotopic position, independent of the others. The assumption of separability underlies many optical imaging experiments and most current thinking about cortical maps. Some forms of inseparability in receptive field structure are known to exist, for example, in the space-time profiles of some simple cell receptive fields (DeAngelis, Ohzawa, & Freeman, 1995). The nonseparability in space and time means that receptive field positions can be considered to shift slightly in space following presentation of a stimulus. However, the shift is small compared with the overall scale of retinotopic mapping and, moreover, affects subunit structure within the envelope, leaving the overall envelope of receptive field sensitivity relatively unchanged. So the retinotopic position represented by the cell's activity is substantially independent of other stimulus values. Although thorough tests do not seem to have been done, there is little or no reported evidence of significant dependence between other pairs of map parameters, including preferred orientation and retinal position, spatial frequency and eye of stimulation, or spatial frequency and orientation.

Recently Basole et al. (2003) have suggested that orientation and direction of motion might not be separable quantities and that this calls into question the idea of separable cortical maps for different features. They used a stimulus consisting of a number of uniformly oriented short bars that were moved in directions that were not orthogonal to the orientation

of the bars. In optically imaged maps of visual cortex, they found a discrepancy between the orientation of the stimulus bars and the predicted regions of cortical response based on the responses to normal-oriented grating stimuli. However, Mante and Carandini (2005) and Baker and Issa (2005) have pointed out that this stimulus spreads spatiotemporal energy into regions of Fourier space whose orientations relative to the origin are different from those of the component bars. In some cases, these orientations can dominate the cortical response. Thus, it is not surprising that for this stimulus, the orientations of the bars do not always predict which regions of cortex will respond. This result can be reconciled with the idea of independent maps of orientation and other features if the term *orientation* is understood to mean "the dominant orientation of particular spatio-temporal frequency components of the image."

6.7 Readout Time. The proposed mechanism involves conduction to and from, and within, a cortical area, integration by the apical dendrite, and two synaptic connections. Thus, single-shot readout should be fast. In the monkey, feedback pathways conduct as rapidly as feedforward pathways, at about 3.5 m/s (Girard et al., 2001). A 1 to 2 cm long pathway would thus take about 3 to 5 ms in each direction. Additional time would be required for lateral spread of activity in the fine tangential branches of feedback projections, which are likely to conduct more slowly, at velocities of 0.1 to 1 m/s. It could take up to 10 ms to propagate activity over branches spreading 5 mm and conducting at around 0.5 m/s, although the exact time would depend on the specific geometry and diameters of the terminal branches. Apical dendritic integration times are about 5 msec (Larkum et al., 1999, 2004; Stuart & Häusser, 2001). Summing and signaling back by the collector neuron might in principle be done in less than 5 ms. This gives a total time of 30 ms or less. It is interesting that experiments using transcranial magnetic stimulation (TMS) show that feedback from V5 to V1 is required for conscious visual perception and that this feedback operates on a comparable timescale of 10 to 40 ms (Pascual-Leone & Walsh, 2001).

These estimates suggest that a significant portion of readout time might be taken up with conduction within the terminal branches of Q-fibers. Minimizing conduction time within these arbors would therefore speed up readout and would have a significant functional advantage. Minimizing conduction time is closely related to the problem of minimizing distances between functionally related groups of neurons, and it has been suggested that many aspects of cortical map organization can be explained as the result of minimizing connection lengths (Durbin & Mitchison, 1990; Chklovskii & Koulakov, 2004). Thus, in addition to minimizing connection lengths and thus cortical volume, an important function of cortical maps may be to minimize readout times and speed up cortical processing.

Acknowledgments

This work was supported by the Canadian Institutes of Health Research and the Natural Sciences and Engineering Research Council of Canada. I thank Graeme Mitchison for comments on an earlier version of the manuscript.

References

Baker, T. I., & Issa, N. P. (2005). Cortical maps of separable tuning properties predict population responses to complex visual stimuli. *J. Neurophysiol.*, *94*, 775–787.

Basole, A., White, L. E., & Fitzpatrick, D. (2003). Mapping multiple features in the population response of visual cortex. *Nature, 423*, 986–990.

Bonhoeffer, T., & Grinvald, A. (1996). Optical imaging based on intrinsic signals: The methodology. In A. W. Toga & J. C. Maziotta (Eds.), *Brain mapping: The methods* (pp. 55–97). San Diego, CA: Academic Press.

Carpenter, G. A., & Grossberg, S. (1987). A massively parallel architecture for a self-organizing neural pattern recognition machine. *Computer Vision, Graphics and Image Processing,* 37, 54–115.

Chawla, D., Rees, G., & Friston, K. J. (1999). The physiological basis of attentional modulation in extrastriate visual areas. *Nature Neuroscience, 2*, 671–676.

Chklovskii, D. B., & Koulakov, A. A. (2004). Maps in the brain: What can we learn from them? *Annu. Rev. Neurosci., 27*, 369–392.

Cohn, T. E., & Lasley, D. J. (1974). Detectability of luminance increment: Effect of spatial uncertainty. *J. Opt. Soc. Am., 64*, 1715–1719.

Dayan, P., Hinton, G. E., Neal, R. M., & Zemel, R. S. (1995). The Helmholtz machine. *Neural Computation, 7*, 889–904.

DeAngelis, G. C., Ohzawa, I., & Freeman, R. D. (1995). Receptive field dynamics in the central visual pathways. *Trends in Neurosciences, 18*, 451–458.

Durbin, R., & Mitchison, G. (1990). A dimension reduction framework for understanding cortical maps. *Nature, 343*, 644–647.

Edwards, A. W. F. (2004). *Cogwheels of the mind*. Baltimore, MD: John Hopkins University Press.

Friston, K. (2005). A theory of cortical responses. *Phil. Trans. R. Soc. Lond. B, 360*, 815–836.

Galuske, R. A. W., Schmidt, K. E., Goebel, R., Lomber, S. G., & Payne, B. R. (2002). The role of feedback in shaping neural representations in cat visual cortex. *Proc. Natl. Acad. Sci., 99*, 17083–17088.

Girard, P., Hupé, J.-M., & Bullier, J. (2001). Feedforward and feedback connections between areas V1 and V2 of the monkey have similar rapid conduction velocities. *J. Neurophysiol., 85*, 1328–1331.

Grossberg, S. (1976). Adaptive pattern classification and universal recoding: II. Feedback, expectation, olfaction, illusions. *Biol. Cybernetics, 23*, 187–202.

Hinton, G., Dayan, P., Frey, B., & Neal, R. (1995). The wake-sleep algorithm for unsupervised neural networks. *Science, 268*, 1158–1161.

Huang, L., Chen, X., & Shou, T. (2004). Spatial frequency-dependent feedback of visual cortical area 21a modulating functional orientation column maps in areas 17 and 18 of the cat. *Brain Res., 998*, 194–201.

Hubel, D. H., & Wiesel, T. N. (1977). Functional architecture of macaque monkey visual cortex. *Proc. R. Soc. Lond.* B, 198, 1–59.

Hübener, M., Shoham, D., Grinvald, A., & Bonhoeffer, T. (1997). Spatial relationships among three columnar systems in cat area 17. *J. Neurosci., 17*, 9270–9284.

Hupé, J.-M., James, A. C., Girard, P., Lomber, S. G., Payne, B., & Bullier, J. (2001). Feedback connections act on the early part of the responses in monkey visual cortex. *J. Neurophysiol., 85*, 134–145.

Hupé, J.-M., James, A. C., Payne, B. R., Lomber, S. G., Girard, P., & Bullier, J. (1998). Cortical feedback improves discrimination between figure and background by V1, V2 and V3 neurons. *Nature, 394*, 784–787.

Kohonen, T. (1982). Self-organized formation of topologically correct feature maps. *Biol. Cybern., 43*, 59–69.

Kosslyn, S. M., & Thompson, W. L. (2003). When is early visual cortex activated during visual mental imagery? *Psychological Bulletin, 129*, 723–746.

Larkum, M. E., Senn, W., & Lüscher, H.-R. (2004). Top-down dendritic input increases the gain of layer 5 pyramidal neurons. *Cerebral Cortex, 14*, 1059–1070.

Larkum, M. E., Zhu, J., & Sakmann, B. (1999). A new cellular mechanism for coupling inputs arriving at different cortical layers. *Nature, 398*, 338–341.

Lee, T. S., & Mumford, D. (2003). Hierarchical Bayesian inference in the visual cortex. *J. Opt. Soc. Am. A, 20*, 1434–1448.

Lindblom, B., & Westheimer, G. (1992). Uncertainty effects in orientation discrimination of foveally seen lines in human observers. *J. Physiol. (London), 454*, 1–8.

Mante, V., & Carandini, M. (2005). Mapping of stimulus energy in visual cortex. *J. Neurophysiol., 94*, 788–798.

Martinez-Trujillo, J. C., & Treue, S. (2004). Feature-based attention increases the selectivity of population responses in primate visual cortex. *Current Biology, 14*, 744–751.

McAdams, C. J., & Maunsell, J. H. (1999). Effects of attention on orientation-tuning functions of single neurons in macaque cortical area V4. *J. Neurosci., 19*, 431–441.

Mountcastle, V. B. (1978). An organizing principle for cerebral function. In G. M. Edelman, & V. B. Mountcastle (Eds.), *The mindful Brain* (pp. 7–50). Cambridge, MA: MIT Press.

Mumford, D. (1992). On the computational architecture of the neo-cortex II. *Biol. Cybern., 66*, 241–251.

Olshausen, B. A., & Field, D. J. (1997). Sparse coding with an overcomplete basis set: A strategy employed by V1? *Vision Res., 37*, 3311–3325.

Pascual-Leone, A., & Walsh, V. (2001). Fast backprojections from the motion to the primary visual area necessary for visual awareness. *Science, 292*, 510–512.

Pouget, A., Dayan, P., & Zemel, R. (2000). Information processing with population codes. *Nature Reviews Neuroscience, 1*, 125–132.

Rao, R., & Ballard, D. (1997). Dynamic model of visual recognition predicts neural response properties in the visual cortex. *Neural Computation, 9*, 721–763.

Reynolds, J. H., Pasternak, T., & Desimone, R. (2000). Attention increases sensitivity of V4 neurons. *Neuron, 26*, 703–714.

Rockland, K. S., & Knutson, T. (2000). Feedback connections from area MT of the squirrel monkey to areas V1 and V2. *J. Comp. Neurol., 425*, 345–368.

Rockland, K. S., & Virga, A. (1989). Terminal arbors of individual "feedback" axons projecting from area V2 to V1 in the macaque monkey: A study using immunohistochemistry of anterogradely transported *Phaseolus vulgaris* leucoagglutinin. *J. Comp. Neurol., 285*, 54–72.

Salin, P.-A., & Bullier, J. (1995). Corticocortical connections in the visual system: Structure and function. *Physiological Reviews, 75*, 107–154.

Schneidmann, E., Berry, M. J., Segev, R., & Bialek, W. (2006). Weak pairwise correlations imply strongly correlated network states in a neural population. *Nature, 440*, 1007–1012.

Shmuel, A., & Grinvald, A. (1996). Functional organization for direction of motion and its relationship to orientation maps in cat area 18. *J. Neurosci., 16*, 6945–6964.

Shmuel, A., Korman, M., Sterkin, A., Harel, M., Ullman, S., Malach, R., & Grinvald, A. (2005). Retinotopic axis specificity and selective clustering of feedback projections from V2 to V1 in the owl monkey. *J. Neurosci., 25*, 2117–2131.

Stuart, G. J., & Häusser, M. (2001). Dendritic coincidence detection of EPSPs and action potentials. *Nature Neuroscience, 4*, 63–71.

Suzuki, W., Saleem, K. S., & Tanaka, K. (2000). Divergent backward projections from the anterior part of the inferotemporal cortex (area TE) in the macaque. *J. Comp. Neurol., 422*, 206–228.

Swindale, N. V. (1991). Coverage and the design of striate cortex. *Biol. Cybern., 65*, 415–424.

Swindale, N. V. (2000). How many maps are there in visual cortex? *Cerebral Cortex, 10*, 633–643.

Swindale, N. V. (2004). How different feature spaces may be represented in cortical maps. *Network, 15*, 217–242.

Swindale, N. V., Shoham, D., Grinvald, A., Bonhoeffer, T., & Hübener, M. (2000). Visual cortex maps are optimized for uniform coverage. *Nature Neurosci., 3*, 822–826.

Treue, S., & Martinez-Trujillo, J. C. (1999). Feature-based attention influences motion processing gain in macaque visual cortex. *Nature, 399*, 575–579.

van der Schaaf, A. & van Hateren, J. H. (1996). Modeling the power spectra of natural images: Statistics and information. *Vision Res., 36*, 2759–2770.

Venn, J. (1880). On the diagrammatic and mechanical representation of propositions and reasonings. *London, Edinburgh and Dublin Philosophical Magazine and Journal of Science [Fifth Series], 9*, 1–18.

Wang, G., Tanifuji, M., & Tanaka, K. (1998). Functional architecture in monkey inferotemporal cortex revealed by in vivo optical imaging. *Neuroscience Research, 32*, 33–46.

Wang, C., Waleszczyk, W. J., Burke, W., & Dreher, B. (2000). Modulatory influences of feedback projections from area 21a on neuronal activities in striate cortex of the cat. *Cerebral Cortex, 10*, 1217–1232.

Watt, R. J. (1987). Scanning from coarse to fine scales in the human visual system after the onset of a stimulus. *J. Opt. Soc. Am., 4A*, 2006–2021.

Weliky, M., Bosking, W. H., & Fitzpatrick, D. A. (1996). Systematic map of direction preference in primary visual cortex. *Nature, 379*, 725–728.

Westheimer, G., & Ley, E. (1996). Temporal uncertainty effects on orientation discrimination and stereoscopic thresholds. *J. Opt. Soc. Am., 13*, 884–886.

Yu, H., Farley, B. J., Jin, D. Z., & Sur, M. (2005). The coordinated mapping of visual space and response features in visual cortex. *Neuron, 47*, 267–280.

Zangaladze, A., Epstein, C. M., Grafton, S. T., & Sathian, K. (1999). Involvement of visual cortex in tactile discrimination of orientation. *Nature, 401*, 587–590.

Received June 26, 2006; accepted December 14, 2006.