2D Observers for Human 3D Object Recognition?

Zili Liu
NEC Research Institute

Daniel Kersten
University of Minnesota

Abstract

Converging evidence has shown that human object recognition depends on familiarity with the images of an object. Further, the greater the similarity between objects, the stronger is the dependence on object appearance, and the more important two-dimensional (2D) image information becomes. These findings, however, do not rule out the use of 3D structural information in recognition, and the degree to which 3D information is used in visual memory is an important issue. Liu, Knill, & Kersten (1995) showed that any model that is restricted to rotations in the image plane of independent 2D templates could not account for human performance in discriminating novel object views. We now present results from models of generalized radial basis functions (GRBF), 2D nearest neighbor matching that allows 2D affine transformations, and a Bayesian statistical estimator that integrates over all possible 2D affine transformations. The performance of the human observers relative to each of the models is better for the novel views than for the familiar template views, suggesting that humans generalize better to novel views from template views. The Bayesian estimator yields the optimal performance with 2D affine transformations and independent 2D templates. Therefore, models of 2D affine matching operations with independent 2D templates are unlikely to account for human recognition performance.

1 Introduction

Object recognition is one of the most important functions in human vision. To understand human object recognition, it is essential to understand how objects are represented in human visual memory. A central component in object recognition is the matching of the stored object representation with that derived from the image input. But the nature of the object representation has to be inferred from recognition performance, by taking into account the contribution from the image information. When evaluating human performance, how can one separate the con-

tributions to performance of the image information from the representation? Ideal observer analysis provides a precise computational tool to answer this question. An ideal observer's recognition performance is restricted only by the available image information and is otherwise optimal, in the sense of statistical decision theory, irrespective of how the model is implemented. A comparison of human to ideal performance (often in terms of efficiency) serves to normalize performance with respect to the image information for the task. We consider the problem of viewpoint dependence in human recognition.

A recent debate in human object recognition has focused on the dependence of recognition performance on viewpoint [1, 6]. Depending on the experimental conditions, an observer's ability to recognize a familiar object from novel viewpoints is impaired to varying degrees. A central assumption in the debate is the equivalence in viewpoint dependence and recognition performance. In other words, the assumption is that viewpoint dependent performance implies a viewpoint dependent representation, and that viewpoint independent performance implies a viewpoint independent representation. However, given that any recognition performance depends on the input image information, which is necessarily viewpoint dependent, the viewpoint dependence of the performance is neither necessary nor sufficient for the viewpoint dependence of the representation. Image information has to be factored out first, and the ideal observer provides the means to do this.

The second aspect of an ideal observer is that it is implementation free. Consider the GRBF model [5], as compared with human object recognition (see below). The model stores a number of 2D templates $\{\mathbf{T}_i\}$ of a 3D object $\mathbf{O}$, and recognizes or rejects a stimulus image $\mathbf{S}$ by the following similarity measure $\Sigma_i c_i \exp\left(\|\mathbf{T}_i - \mathbf{S}\|^2/2\sigma^2\right)$, where c_i and σ are constants. The model's performance as a function of viewpoint parallels that of human observers. This observation has led to the conclusion that the human visual system may indeed, as does the model, use 2D stored views with GRBF interpolation to recognize 3D objects [2]. Such a conclusion, however, overlooks implementational constraints in the model, because the model's performance also depends on its implementations. Conceivably, a model with some 3D information of the objects can also mimic human performance, so long as it is appropriately implemented. There are typically too many possible models that can produce the same pattern of results.

In contrast, an ideal observer computes the optimal performance that is only limited by the stimulus information and the task. We can define constrained ideals that are also limited by explicitly specified assumptions (e.g., a class of matching operations). Such a model observer therefore yields the best possible performance among the class of models with the same stimulus input and assumptions. In this paper, we are particularly interested in constrained ideal observers that are restricted in functionally significant aspects (e.g., a 2D ideal observer that stores independent 2D templates and has access only to 2D affine transformations). The key idea is that a constrained ideal observer is the best in its class. So if humans outperform this ideal observer, they must have used more than what is available to the ideal. The conclusion that follows is strong: not only does the constrained ideal fail to account for human performance, but the whole class of its implementations are also falsified.

A crucial question in object recognition is the extent to which human observers model the geometric variation in images due to the projection of a 3D object onto a 2D image. At one extreme, we have shown that any model that compares the image to independent views (even if we allow for 2D rigid transformations of the input image) is insufficient to account for human performance . At the other extreme, it is unlikely that variation is modeled in terms of rigid transformation of a 3D object

template in memory. A possible intermediate solution is to match the input image to stored views, subject to 2D affine deformations. This is reasonable because 2D affine transformations approximate 3D variation over a limited range of viewpoint change.

In this study, we test whether any model limited to the independent comparison of 2D views, but with 2D affine flexibility, is sufficient to account for viewpoint dependence in human recognition. In the following section, we first define our experimental task, in which the computational models yield the provably best possible performance under their specified conditions. We then review the 2D ideal observer and GRBF model derived in [4], and the 2D affine nearest neighbor model in [8]. Our principal theoretical result is a closed-form solution of a Bayesian 2D affine ideal observer. We then compare human performance with the 2D affine ideal model, as well as the other three models. In particular, if humans can classify novel views of an object better than the 2D affine ideal, then our human observers must have used more information than that embodied by that ideal.

2 The observers

Let us first define the task. An observer looks at the 2D images of a 3D wire frame object from a number of viewpoints. These images will be called templates $\{\mathbf{T}_i\}$. Then two distorted copies of the original 3D object are displayed. They are obtained by adding 3D Gaussian positional noise (i.i.d.) to the vertices of the original object. One distorted object is called the target, whose Gaussian noise has a constant variance. The other is the distractor, whose noise has a larger variance that can be adjusted to achieve a criterion level of performance. The two objects are displayed from the same viewpoint in parallel projection, which is either from one of the template views, or a novel view due to 3D rotation. The task is to choose the one that is more similar to the original object. The observer's performance is measured by the variance (threshold) that gives rise to 75% correct performance. The optimal strategy is to choose the stimulus $\mathbf{S}$ with a larger probability $p(\mathbf{O}|\mathbf{S})$. From Bayes' rule, this is to choose the larger of $p(\mathbf{S}|\mathbf{O})$.

Assume that the models are restricted to 2D transformations of the image, and cannot reconstruct the 3D structure of the object from its independent templates $\{\mathbf{T}_i\}$. Assume also that the prior probability $p(\mathbf{T}_i)$ is constant. Let us represent $\mathbf{S}$ and $\mathbf{T}_i$ by their (x, y) vertex coordinates: $(\ \mathbf{X}\quad \mathbf{Y}\ )^T$, where $\mathbf{X} = (x^1, x^2, \ldots, x^n)$, $\mathbf{Y} = (y^1, y^2, \ldots, y^n)$. We assume that the correspondence between $\mathbf{S}$ and $\mathbf{T}_i$ is solved up to a reflection ambiguity, which is equivalent to an additional template: $\mathbf{T}_i^r = (\ \mathbf{X}^r\quad \mathbf{Y}^r\ )^T$, where $\mathbf{X}^r = (x^n, \ldots, x^2, x^1)$, $\mathbf{Y}^r = (y^n, \ldots, y^2, y^1)$. We still denote the template set as $\{\mathbf{T}_i\}$. Therefore,

$$p(\mathbf{S}|\mathbf{O}) = \Sigma p(\mathbf{S}|\mathbf{T}_i)p(\mathbf{T}_i). \tag{1}$$

In what follows, we will compute $p(\mathbf{S}|\mathbf{T}_i)p(\mathbf{T}_i)$, with the assumption that $\mathbf{S} = \mathcal{F}(\mathbf{T}_i) + \mathbf{N}(\mathbf{0}, \sigma\mathbf{I}_{2n})$, where $\mathbf{N}$ is the Gaussian distribution, $\mathbf{I}_{2n}$ the $2n \times 2n$ identity matrix, and $\mathcal{F}$ a 2D transformation. For the 2D ideal observer, $\mathcal{F}$ is a rigid 2D rotation. For the GRBF model, $\mathcal{F}$ assigns a linear coefficient to each template $\mathbf{T}_i$, in addition to a 2D rotation. For the 2D affine nearest neighbor model, $\mathcal{F}$ represents the 2D affine transformation that minimizes $\|\mathbf{S} - \mathbf{T}_i\|^2$, after $\mathbf{S}$ and $\mathbf{T}_i$ are normalized in size. For the 2D affine ideal observer, $\mathcal{F}$ represents all possible 2D affine transformations applicable to $\mathbf{T}_i$.

2.1 The 2D ideal observer

The templates are the original 2D images, their mirror reflections, and 2D rotations (in angle ϕ) in the image plane. Assume that the stimulus S is generated by adding Gaussian noise to a template, the probability $p(\mathbf{S}|\mathbf{O})$ is an integration over all templates and their reflections and rotations. The detailed derivation for the 2D ideal and the GRBF model can be found in [4].

$$\Sigma p(\mathbf{S}|\mathbf{T}_i)p(\mathbf{T}_i) \propto \Sigma \int d\phi \exp\left(-\|\mathbf{S}-\mathbf{T}_i(\phi)\|^2/2\sigma^2\right). \tag{2}$$

2.2 The GRBF model

The model has the same template set as the 2D ideal observer does. Its training requires that $\Sigma_i \int_0^{2\pi} d\phi c_i(\phi) N(\|\mathbf{T}_j - \mathbf{T}_i(\phi)\|, \sigma) = 1$, $j = 1, 2, \ldots$, with which $\{c_i\}$ can be obtained optimally using singular value decomposition. When a pair of new stimuli $\{\mathbf{S}\}$ are presented, the optimal decision is to choose the one that is closer to the learned prototype, in other words, the one with a smaller value of

$$\|1 - \Sigma \int_0^{2\pi} d\phi c_i(\phi) \exp\left(-\frac{\|\mathbf{S}-\mathbf{T}_i(\phi)\|^2}{2\sigma^2}\right)\|. \tag{3}$$

2.3 The 2D affine nearest neighbor model

It has been proved in [8] that the smallest Euclidean distance $D(\mathbf{S}, \mathbf{T})$ between $\mathbf{S}$ and $\mathbf{T}$ is, when $\mathbf{T}$ is allowed a 2D affine transformation, $\mathbf{S} \to \mathbf{S}/\|\mathbf{S}\|$, $\mathbf{T} \to \mathbf{T}/\|\mathbf{T}\|$,

$$D^2(\mathbf{S}, \mathbf{T}) = 1 - tr(\mathbf{S}^+\mathbf{S} \cdot \mathbf{T}^T\mathbf{T})/\|\mathbf{T}\|^2, \tag{4}$$

where *tr* strands for *trace*, and $\mathbf{S}^+ = \mathbf{S}^T(\mathbf{S}\mathbf{S}^T)^{-1}$. The optimal strategy, therefore, is to choose the $\mathbf{S}$ that gives rise to the larger of $\Sigma \exp\left(-D^2(\mathbf{S}, \mathbf{T}_i)/2\sigma^2\right)$, or the smaller of $\Sigma D^2(\mathbf{S}, \mathbf{T}_i)$. (Since no probability is defined in this model, both measures will be used and the results from the better one will be reported.)

2.4 The 2D affine ideal observer

We now calculate the Bayesian probability by assuming that the prior probability distribution of the 2D affine transformation, which is applied to the template $\mathbf{T}_i$, $\mathbf{AT} + \mathbf{T_r} = \begin{pmatrix} a & b \\ c & d \end{pmatrix} \mathbf{T}_i + \begin{pmatrix} t_x & \cdots & t_x \\ t_y & \cdots & t_y \end{pmatrix}$, obeys a Gaussian distribution $\mathbf{N}(\mathbf{X_0}, \gamma\mathbf{I_6})$, where $\mathbf{X_0}$ is the identity transformation $\mathbf{X}_0^T = (a, b, c, d, t_x, t_y) = (1, 0, 0, 1, 0, 0)$. We have

$$\begin{aligned} \Sigma p(\mathbf{S}|\mathbf{T}_i) &= \Sigma \int_{-\infty}^{\infty} d\mathbf{X} \exp\left(-\|\mathbf{AT}_i + \mathbf{T_r} - \mathbf{S}\|^2/2\sigma^2\right) &(5)\\ &= \Sigma C(n, \sigma, \gamma) det^{-1}\left(\mathbf{Q}_i'\right) \exp\left(tr\left(\mathbf{K}_i^T\mathbf{Q}_i(\mathbf{Q}_i')^{-1}\mathbf{Q}_i\mathbf{K}_i\right)/2\sigma^2\right), &(6)\end{aligned}$$

where $C(n, \sigma, \gamma)$ is a function of n, σ, γ; $\mathbf{Q}' = \mathbf{Q} + \gamma^{-2}\mathbf{I}_2$, and

$$\mathbf{Q} = \begin{pmatrix} \mathbf{X_T} \cdot \mathbf{X_T} & \mathbf{X_T} \cdot \mathbf{Y_T} \\ \mathbf{Y_T} \cdot \mathbf{X_T} & \mathbf{Y_T} \cdot \mathbf{Y_T} \end{pmatrix}, \mathbf{QK} = \begin{pmatrix} \mathbf{X_T} \cdot \mathbf{X_S} & \mathbf{Y_T} \cdot \mathbf{X_S} \\ \mathbf{X_T} \cdot \mathbf{Y_S} & \mathbf{Y_T} \cdot \mathbf{Y_S} \end{pmatrix} + \gamma^{-2}\mathbf{I}_2. \tag{7}$$

The free parameters are γ and the number of 2D rotated copies for each $\mathbf{T}_i$ (since a 2D affine transformation implicitly includes 2D rotations, and since a specific prior probability distribution $\mathbf{N}(\mathbf{X_0}, \gamma\mathbf{I})$ is assumed, both free parameters should be explored together to search for the optimal results).

Figure 1: Stimulus classes with increasing structural regularity: Balls, Irregular, Symmetric, and V-Shaped. There were three objects in each class in the experiment.

2.5 The human observers

Three naive subjects were tested with four classes of objects: Balls, Irregular, Symmetric, and V-Shaped (Fig. 1). There were three objects in each class. For each object, 11 template views were learned by rotating the object 60°/step, around the X- and Y-axis, respectively. The 2D images were generated by orthographic projection, and viewed monocularly. The viewing distance was 1.5 m. During the test, the standard deviation of the Gaussian noise added to the target object was $\sigma_t = 0.254$ cm. No feedback was provided.

Because the image information available to the humans was more than what was available to the models (shading and occlusion in addition to the (x, y) positions of the vertices), both learned and novel views were tested in a randomly interleaved fashion. Therefore, the strategy that humans used in the task for the learned and novel views should be the same. The number of self-occlusions, which in principle provided relative depth information, was counted and was about equal in both learned and novel view conditions. The shading information was also likely to be equal for the learned and novel views. Therefore, this additional information was about equal for the learned and novel views, and should not affect the comparison of the performance (humans relative to a model) between learned and novel views. We predict that if the humans used a 2D affine strategy, then their performance relative to the 2D affine ideal observer should not be higher for the novel views than for the learned views. One reason to use the four classes of objects with increasing structural regularity is that structural regularity is a 3D property (e.g., 3D Symmetric vs. Irregular), which the 2D models cannot capture. The exception is the planar V-Shaped objects, for which the 2D affine models completely capture 3D rotations, and are therefore the "correct" models. The V-Shaped objects were used in the 2D affine case as a benchmark. If human performance increases with increasing structural regularity of the objects, this would lend support to the hypothesis that humans have used 3D information in the task.

2.6 Measuring performance

A stair-case procedure [7] was used to track the observers' performance at 75% correct level for the learned and novel views, respectively. There were 120 trials for the humans, and 2000 trials for each of the models. For the GRBF model, the standard deviation of the Gaussian function was also sampled to search for the best result for the novel views for each of the 12 objects, and the result for the learned views was obtained accordingly. This resulted in a conservative test of the hypothesis of a GRBF model for human vision for the following reasons: (1) Since no feedback was provided in the human experiment and the learned and novel views were randomly intermixed, it is not straightforward for the model to find the best standard deviation for the novel views, particularly because the best standard deviation for the novel views was not the same as that for the learned

ones. The performance for the novel views is therefore the upper limit of the model's performance. (2) The subjects' performance relative to the model will be defined as statistical efficiency (see below). The above method will yield the lowest possible efficiency for the novel views, and a higher efficiency for the learned views, since the best standard deviation for the novel views is different from that for the learned views. Because our hypothesis depends on a higher statistical efficiency for the novel views than for the learned views, this method will make such a putative difference even smaller. Likewise, for the 2D affine ideal, the number of 2D rotated copies of each template $\mathbf{T}_i$ and the value γ were both extensively sampled, and the best performance for the novel views was selected accordingly. The result for the learned views corresponding to the same parameters was selected. This choice also makes it a conservative hypothesis test.

3 Results

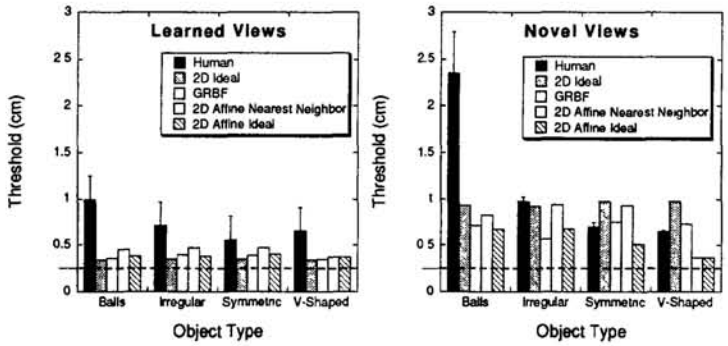

Figure 2: The threshold standard deviation of the Gaussian noise, added to the distractor in the test pair, that keeps an observer's performance at the 75% correct level, for the learned and novel views, respectively. The dotted line is the standard deviation of the Gaussian noise added to the target in the test pair.

Fig. 2 shows the threshold performance. We use statistical efficiency $\mathcal{E}$ to compare human to model performance. $\mathcal{E}$ is defined as the information used by humans relative to the ideal observer [3] : $\mathcal{E} = (d'_{human}/d'_{ideal})^2$, where d' is the discrimination index. We have shown in [4] that, in our task, $\mathcal{E} = \left((\sigma^{ideal}_{distractor})^2 - (\sigma_{target})^2\right) / \left((\sigma^{human}_{distractor})^2 - (\sigma_{target})^2\right)$, where σ is the threshold. Fig. 3 shows the statistical efficiency of the human observers relative to each of the four models.

We note in Fig. 3 that the efficiency for the novel views is higher than those for the learned views (several of them even exceeded 100%), except for the planar V-Shaped objects. We are particularly interested in the Irregular and Symmetric objects in the 2D affine ideal case, in which the pairwise comparison between the learned and novel views across the six objects and three observers yielded a significant difference (binomial, $p < 0.05$). This suggests that the 2D affine ideal observer cannot account for the human performance, because if the humans used a 2D affine template matching strategy, their relative performance for the novel views cannot be better than for the learned views. We suggest therefore that 3D information was used by the human observers (e.g., 3D symmetry). This is supported in addition by the increasing efficiencies as the structural regularity increased from the Balls, Irregular, to Symmetric objects (except for the V-Shaped objects with 2D affine models).

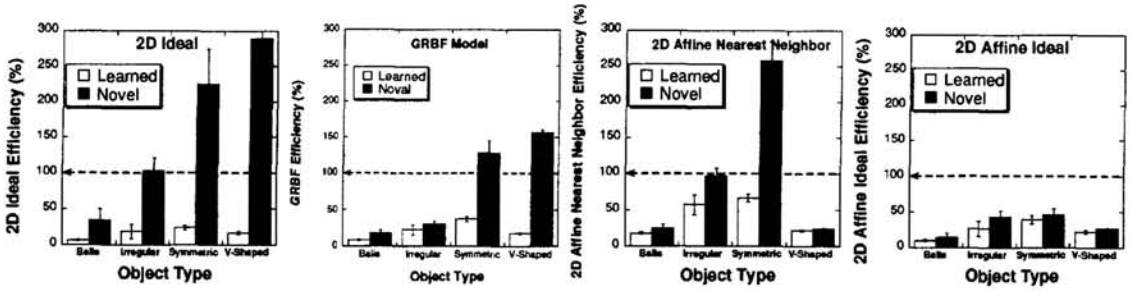

Figure 3: Statistical efficiencies of human observers relative to the 2D ideal observer, the GRBF model, the 2D affine nearest neighbor model, and the 2D affine ideal observer.

4 Conclusions

Computational models of visual cognition are subject to information theoretic as well as implementational constraints. When a model's performance mimics that of human observers, it is difficult to interpret which aspects of the model characterize the human visual system. For example, human object recognition could be simulated by both a GRBF model and a model with partial 3D information of the object. The approach we advocate here is that, instead of trying to mimic human performance by a computational model, one designs an implementation-free model for a specific recognition task that yields the best possible performance under explicitly specified computational constraints. This model provides a well-defined benchmark for performance, and if human observers outperform it, we can conclude firmly that the humans must have used better computational strategies than the model. We showed that models of independent 2D templates with 2D linear operations cannot account for human performance. This suggests that our human observers may have used the templates to reconstruct a representation of the object with some (possibly crude) 3D structural information.

References

[1] Biederman I and Gerhardstein P C. Viewpoint dependent mechanisms in visual object recognition: a critical analysis. J. Exp. Psych.: HPP, 21:1506–1514, 1995.

[2] Bülthoff H H and Edelman S. Psychophysical support for a 2D view interpolation theory of object recognition. Proc. Natl. Acad. Sci., 89:60–64, 1992.

[3] Fisher R A. Statistical Methods for Research Workers. Oliver and Boyd, Edinburgh, 1925.

[4] Liu Z, Knill D C, and Kersten D. Object classification for human and ideal observers. Vision Research, 35:549–568, 1995.

[5] Poggio T and Edelman S. A network that learns to recognize three-dimensional objects. Nature, 343:263–266, 1990.

[6] Tarr M J and Bülthoff H H. Is human object recognition better described by geon-structural-descriptions or by multiple-views? J. Exp. Psych.: HPP, 21:1494–1505, 1995.

[7] Watson A B and Pelli D G. QUEST: A Bayesian adaptive psychometric method. Perception and Psychophysics, 33:113–120, 1983.

[8] Werman M and Weinshall D. Similarity and affine invariant distances between 2D point sets. IEEE PAMI, 17:810–814, 1995.

Toward a Single-Cell Account for Binocular Disparity Tuning: An Energy Model May be Hiding in Your Dendrites

Bartlett W. Mel
Department of Biomedical Engineering
University of Southern California, MC 1451
Los Angeles, CA 90089
mel@quake.usc.edu

Daniel L. Ruderman
The Salk Institute
10010 N. Torrey Pines Road
La Jolla, CA 92037
ruderman@salk.edu

Kevin A. Archie
Neuroscience Program
University of Southern California
Los Angeles, CA 90089
karchie@quake.usc.edu

Abstract

Hubel and Wiesel (1962) proposed that complex cells in visual cortex are driven by a pool of simple cells with the same preferred orientation but different spatial phases. However, a wide variety of experimental results over the past two decades have challenged the pure hierarchical model, primarily by demonstrating that many complex cells receive monosynaptic input from unoriented LGN cells, or do not depend on simple cell input. We recently showed using a detailed biophysical model that nonlinear interactions among synaptic inputs to an excitable dendritic tree could provide the nonlinear subunit computations that underlie complex cell responses (Mel, Ruderman, & Archie, 1997). This work extends the result to the case of complex cell binocular disparity tuning, by demonstrating in an isolated model pyramidal cell (1) disparity tuning at a resolution much finer than the the overall dimensions of the cell's receptive field, and (2) systematically shifted optimal disparity values for rivalrous pairs of light and dark bars—both in good agreement with published reports (Ohzawa, DeAngelis, & Freeman, 1997). Our results reemphasize the potential importance of intradendritic computation for binocular visual processing in particular, and for cortical neurophysiology in general.

1 Introduction

Binocular disparity is a powerful cue for depth in vision. The neurophysiological basis for binocular disparity processing has been of interest for decades, spawned by the early studies of Hubel and Wiesel (1962) showing neurons in primary visual cortex which could be driven by both eyes. Early qualitative models for disparity tuning held that a binocularly driven neuron could represent a particular disparity (zero, near, or far) via a relative shift of receptive field (RF) centers in the right and left eyes. According to this model, a binocular cell fires maximally when an optimal stimulus, e.g. an edge of a particular orientation, is simultaneously centered in the left and right eye receptive fields, corresponding to a stimulus at a specific depth relative to the fixation point. An account of this kind is most relevant to the case of a cortical "simple" cell, whose phase-sensitivity enforces a preference for a particular absolute location and contrast polarity of a stimulus within its monocular receptive fields.

This global receptive field shift account leads to a conceptual puzzle, however, when binocular *complex* cell receptive fields are considered instead, since a complex cell can respond to an oriented feature nearly independent of position within its monocular receptive field. Since complex cell receptive field diameters in the cat lie in the range of 1-3 degrees, the excessive "play" in their monocular receptive fields would seem to render complex cells incapable of signaling disparity on the much finer scale needed for depth perception (measured in minutes).

Intriguingly, various authors have reported that a substantial fraction of complex cells in cat visual cortex are in fact tuned to left-right disparities much finer than that suggested by the size of the monocular RF's. For such cells, a stimulus delivered at the proper disparity, regardless of absolute position in either eye, produces a neural response in excess of that predicted by the sum of the monocular responses (Pettigrew, Nikara, & Bishop, 1968; Ohzawa, DeAngelis, & Freeman, 1990; Ohzawa et al., 1997). Binocular responses of this type suggest that for these cells, the left and right RF's are combined via a correlation operation rather than a simple sum (Nishihara & Poggio, 1984; Koch & Poggio, 1987). This computation has also been formalized in terms of an "energy" model (Ohzawa et al., 1990, 1997), building on the earlier use of energy models to account for complex cell orientation tuning (Pollen & Ronner, 1983) and direction selectivity (Adelson & Bergen, 1985). In an energy model for binocular disparity tuning, sums of linear Gabor filter outputs representing left and right receptive fields are squared to produce the crucial multiplicative cross terms (Ohzawa et al., 1990, 1997).

Our previous biophysical modeling work has shown that the dendritic tree of a cortical pyramidal cells is well suited to support an approximative high-dimensional quadratic input-output relation, where the second-order multiplicative cross terms arise from local interactions among synaptic inputs carried out in quasi-isolated dendritic "subunits" (Mel, 1992b, 1992a, 1993). We recently applied these ideas to show that the position-invariant orientation tuning of a monocular complex cell could be computed within the dendrites of a single cortical cell, based exclusively upon excitatory inputs from a uniform, overlapping population of unoriented ON and OFF cells (Mel et al., 1997). Given the similarity of the "energy" formulations previously proposed to account for orientation tuning and binocular disparity tuning, we hypothesized that a similar type of dendritic subunit computation could underlie disparity tuning in a binocularly driven complex cell.

Parameter	Value
R_m	$10\text{k}\Omega\text{cm}^2$
R_a	$200\Omega\text{cm}$
C_m	$1.0\mu\text{F/cm}^2$
V_{rest}	-70 mV
Compartments	615
Somatic $\bar{g}_{\text{Na}}$, $\bar{g}_{\text{DR}}$	$0.20, 0.12\ \text{S/cm}^2$
Dendritic $\bar{g}_{\text{Na}}$, $\bar{g}_{\text{DR}}$	$0.05, 0.03\ \text{S/cm}^2$
Input frequency	0 – 100 Hz
$\bar{g}_{\text{AMPA}}$	0.027 nS – 0.295 nS
$\tau_{\text{AMPA}}(on, off)$	0.5 ms, 3 ms
$\bar{g}_{\text{NMDA}}$	0.27 nS – 2.95 nS
$\tau_{\text{NMDA}}(on, off)$	0.5 ms, 50 ms
E_{syn}	0 mV

Table 1: Biophysical simulation parameters. Details of HH channel implementation are given elsewhere (Mel, 1993); original HH channel implementation courtesy Ojvind Bernander and Rodney Douglas. In order that local EPSP size be held approximately constant across the dendritic arbor, peak synaptic conductance at dendritic location x was approximately scaled to the local input resistance (inversely), given by $\bar{g}_{\text{syn}}(x) = c/\tilde{R}_{in}(x)$, where c was a constant, and $\tilde{R}_{in}(x) = \max(R_{in}(x), 200M\Omega)$. Input resistance $R_{\text{in}}(x)$ was measured for a passive cell. Thus $\bar{g}_{\text{syn}}$ was identical for all dendritic sites with input resistance below $200M\Omega$, and was given by the larger conductance value shown; roughly 50% of the tree fell within a factor of 2 of this value. Peak conductances at the finest distal tips were smaller by roughly a factor of 10 (smaller number shown). Somatic input resistance was near $24M\Omega$. The peak synaptic conductance values used were such that the ratio of steady state current injection through NMDA vs. AMPA channels was 1.2 ± 0.4. Both AMPA and NMDA-type synaptic conductances were modeled using the kinetic scheme of Destexhe et al. (1994); synaptic activation and inactivation time constants are shown for each.

2 Methods

Compartmental simulations of a pyramidal cell from cat visual cortex (morphology courtesy of Rodney Douglas and Kevan Martin) were carried out in NEURON (Hines, 1989); simulation parameters are summarized in Table 1. The soma and dendritic membrane contained Hodgkin-Huxley-type (HH) voltage-dependent sodium and potassium channels. Following evidence for higher spike thresholds and decremental propagation in dendrites (Stuart & Sakmann, 1994), HH channel density was set to a uniform, 4-fold lower value in the dendritic membrane relative to that of the cell body. Excitatory synapses from LGN cells included both NMDA and AMPA-type synaptic conductances. Since the cell was considered to be isolated from the cortical network, inhibitory input was not modeled. Cortical cell responses were reported as average spike rate recorded at the cell body over the 500 ms stimulus period, excluding the 50 ms initial transient.

The binocular LGN consisted of two copies of the monocular LGN model used previously (Mel et al., 1997), each consisting of a superimposed pair of 64x64 ON and OFF subfields. LGN cells were modeled as linear, half-rectified center-surround filters with centers 7 pixels in width. We randomly subsampled the left and right LGN arrays by a factor of 16 to yield 1,024 total LGN inputs to the pyramidal cell.

A developmental principle was used to determine the spatial arrangement of these 1,024 synaptic contacts onto the dendritic branches of the cortical cell, as follows. A virtual stimulus ensemble was defined for the cell, consisting of the complete set of single vertical light or dark bars presented binocularly at zero-disparity within the cell's receptive field. Within this ensemble, strong pairwise correlations existed among cells falling into vertically aligned groups of the same (ON or OFF) type, and cells in the vertical column at zero horizontal disparity in the other eye. These binocular cohorts of highly correlated LGN cells were labeled mutual "friends". Progressing through the dendritic tree in depth first order, a randomly chosen LGN cell was assigned to the first dendritic site. A randomly chosen "friend" of hers was assigned to the second site, the third site was assigned to a friend of the site 2 input, etc., until all friends in the available subsample were assigned (4 from each eye, on average). If the friends of the connection at site i were exhausted, a new LGN cell was chosen at random for site $i+1$. In earlier work, this type of synaptic arrangement was shown to be the outcome of a Hebb-type correlational learning rule, in which random, activity independent formation of synaptic contacts acted to slowly randomize the axo-dendritic interface, shaped by Hebbian stabilization of synaptic contacts based on their short-range correlations with other synapses.

3 Results

Model pyramidal cells configured in this way exhibited prominent phase-invariant orientation tuning, the hallmark response property of the visual complex cell. Multiple orientation tuning curves are shown, for example, for a monocular complex cell, giving rise to strong tuning for light and dark bars across the receptive field (fig. 1). The bold curve shows the average of all tuning curves for this cell; the half-width at half max is 25°, in the normal range for complex cells in cat visual cortex (Orban, 1984). When the spatial arrangement of LGN synaptic contacts onto the pyramidal cell dendrites was randomly scrambled, leaving all other model parameters unchanged, orientation tuning was abolished in this cell (right frame), confirming the crucial role of spatially-mediated nonlinear synaptic interactions (average curve from left frame is reproduced for comparison).

Disparity-tuning in an orientation-tuned binocular model cell is shown in fig. 2, compared to data from a complex cell in cat visual cortex (adapted from Ohzawa et al. (1997)). Responses to contrast matched (light-light) and contrast non-matched (light-dark) bar pairs were subtracted to produce these plots. The strong diagonal structure indicates that both the model and real cells responded most vigorously when contrast-matched bars were presented at the same horizontal position in the left and right-eye RF's (i.e. at zero-disparity), whereas peak responses to contrast-non-matched bars occured at symmetric near and far, non-zero disparities.

4 Discussion

The response pattern illustrated in fig. 2A is highly similar to the response generated by an analytical binocular energy model for a complex cell (Ohzawa et al., 1997):

$$\begin{aligned} R_C(X_L, X_R) &= \{\exp(-kX_L^2)\cos(2\pi f X_L) + \exp(-kX_R^2)\cos(2\pi f X_R)\}^2 + \\ &\quad \{\exp(-kX_L^2)\sin(2\pi f X_L) + \exp(-kX_R^2)\sin(2\pi f X_R)\}^2, \end{aligned} \tag{1}$$

where X_L and X_R are the horizontal bar positions to the two eyes, k is the factor

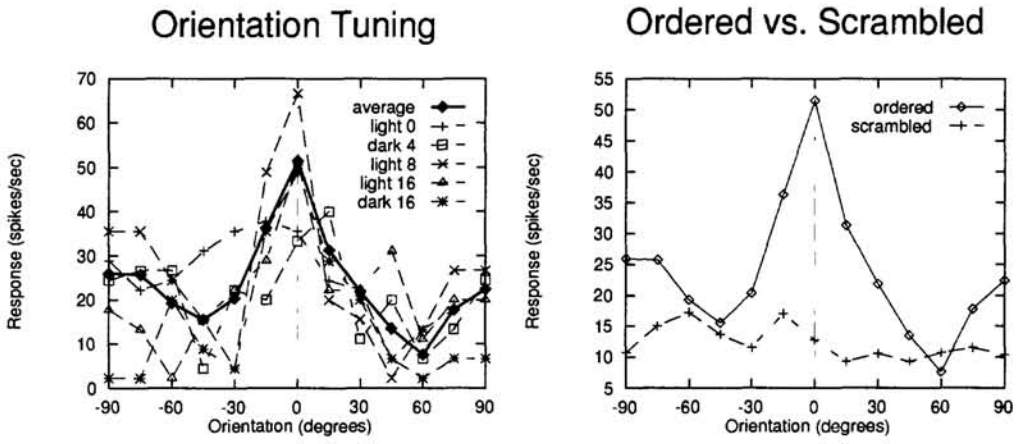

Figure 1: Orientation tuning curves are shown in the left frame for light and dark bars at 3 arbitrary positions. Essentially similar responses were seen at other receptive field positions, and for other complex cells. Bold trace indicates average of tuning curves at positions 0, 1, 2, 4, 8, and 16 for light and dark bars. Similar form of 6 curves shown reflects the translation-invariance of the cell's response to oriented stimuli, and symmetry with respect to ON and OFF input. Orientation tuning is eliminated when the spatial arrangement of LGN synapses onto the model cell dendrites is randomly scrambled (right frame).

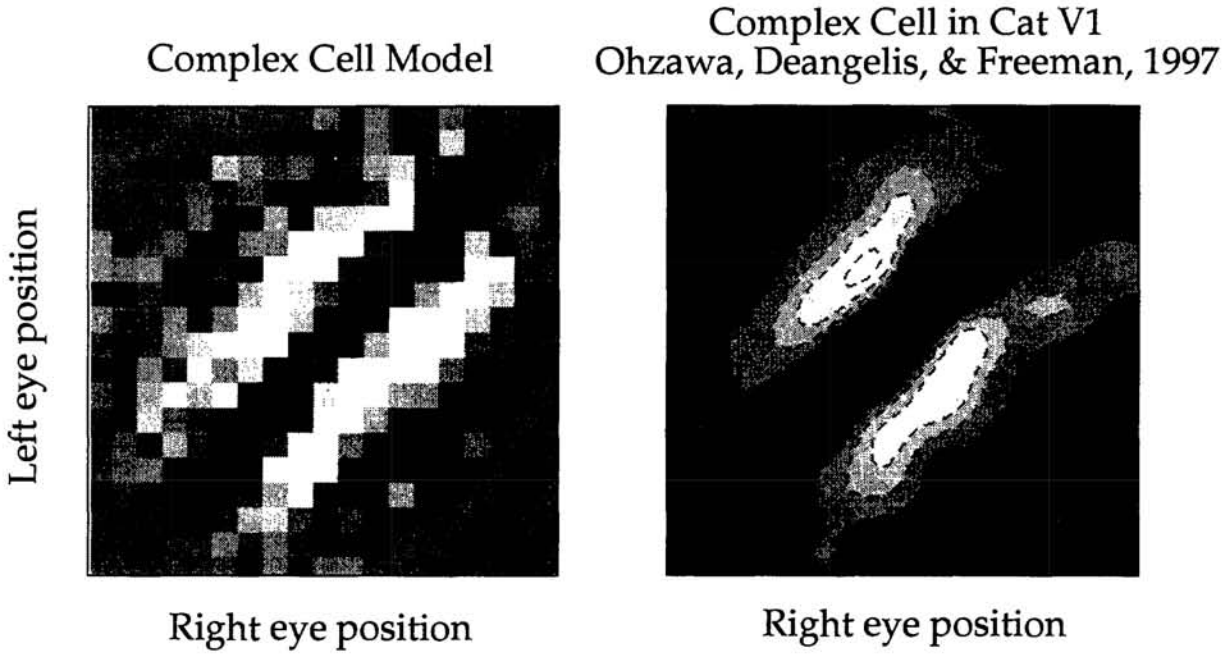

Figure 2: Comparison of disparity tuning in model complex cell to that of a binocular complex cell from cat visual cortex. Light or dark bars were presented simultaneously to the left and right eyes. Bars could be of same polarity in both eyes (light, light) or different polarity (light, dark); cell responses for these two cases were subtracted to produce plot shown in left frame. Right frame shows data similarly displayed for a binocular complex cell in cat visual cortex (adapted from Ohzawa et al. (1997)).

that determines the width of the subunit RF's, and f is the spatial frequency.

In lieu of literal simple cell "subunits", the present results indicate that the subunit computations associated with the terms of an energy model could derive largely from synaptic interactions within the dendrites of the individual cortical cell, driven exclusively by excitatory inputs from unoriented, monocular ON and OFF cells drawn from a uniform overlapping spatial distribution. While lateral inhibition and excitation play numerous important roles in cortical computation, the present results suggest they are not essential for the basic features of the nonlinear disparity tuned responses of cortical complex cells. Further, these results address the paradox as to how inputs from both unoriented LGN cells and oriented simple cells can coexist without conflict within the dendrites of a single complex cell.

A number of controls from previous work suggest that this type of subunit processing is very robustly computed in the dendrites of an individual neuron, with little sensitivity to biophysical parameters and modeling assumptions, including details of the algorithm used to spatially organize the geniculo-cortical projection, specifics of cell morphology, synaptic activation density across the dendritic tree, passive membrane and cytoplasmic parameters, and details of the kinetics, voltage-dependence, or spatial distribution of the voltage-dependent dendritic channels.

One important difference between a standard energy model and the intradendritic responses generated in the present simulation experiments is that the energy model has oriented RF structure at the linear (simple-cell-like) stage, giving rise to oriented, antagonistic ON-OFF subregions (Movshon, Thompson, & Tolhurst, 1978), whereas the linear stage in our model gives rise to center-surround antagonism only within individual LGN receptive fields. Put another way, the LGN-derived subunits in the present model cannot provide all the negative cross-terms that appear in the energy model equations, specifically for pairs of pixels that fall outside the range of a single LGN receptive field.

While the present simulations involve numerous simplifications relative to the full complexity of the cortical microcircuit, the results nonetheless emphasize the potential importance of intradendritic computation in visual cortex.

Acknowledgements

Thanks to Ken Miller, Allan Dobbins, and Christof Koch for many helpful comments on this work. This work was funded by the National Science Foundation and the Office of Naval Research, and by a Sloan Foundation Fellowship (D.R.).

References

Adelson, E., & Bergen, J. (1985). Spatiotemporal energy models for the perception of motion. *J. Opt. Soc. Amer., A 2*, 284–299.

Hines, M. (1989). A program for simulation of nerve equations with branching geometries. *Int. J. Biomed. Comput., 24*, 55–68.

Hubel, D., & Wiesel, T. (1962). Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. *J. Physiol., 160*, 106–154.

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

Mel, B. (1992a). The clusteron: Toward a simple abstraction for a complex neuron. In Moody, J., Hanson, S., & Lippmann, R. (Eds.), *Advances in Neural*

Information Processing Systems, vol. 4, pp. 35–42. Morgan Kaufmann, San Mateo, CA.

Mel, B. (1992b). NMDA-based pattern discrimination in a modeled cortical neuron. *Neural Computation, 4*, 502–516.

Mel, B. (1993). Synaptic integration in an excitable dendritic tree. *J. Neurophysiol., 70*(3), 1086–1101.

Mel, B., Ruderman, D., & Archie, K. (1997). Complex-cell responses derived from center-surround inputs: the surprising power of intradendritic computation. In Mozer, M., Jordan, M., & Petsche, T. (Eds.), *Advances in Neural Information Processing Systems*, Vol. 9, pp. 83–89. MIT Press, Cambridge, MA.

Movshon, J., Thompson, I., & Tolhurst, D. (1978). Receptive field organization of complex cells in the cat's striate cortex. *J. Physiol., 283*, 79–99.

Nishihara, H., & Poggio, T. (1984). Stereo vision for robotics. In Brady, & Paul (Eds.), *Proceedings of the First International Symposium of Robotics Research*, pp. 489–505. MIT Press, Cambridge, MA.

Ohzawa, I., DeAngelis, G., & Freeman, R. (1990). Stereoscopic depth discrimination in the visual cortex: Neurons ideally suited as disparity detectors. *Science, 249*, 1037–1041.

Ohzawa, I., DeAngelis, G., & Freeman, R. (1997). Encoding of binocular disparity by complex cells in the cat's visual cortex. *J. Neurophysiol., June.*

Orban, G. (1984). *Neuronal operations in the visual cortex.* Springer Verlag, New York.

Pettigrew, J., Nikara, T., & Bishop, P. (1968). Responses to moving slits by single units in cat striate cortex. *Exp. Brain Res., 6*, 373–390.

Pollen, D., & Ronner, S. (1983). Visual cortical neurons as localized spatial frequency filters. *IEEE Trans. Sys. Man Cybern., 13*, 907–916.

Stuart, G., & Sakmann, B. (1994). Active propagation of somatic action potentials into neocortical pyramidal cell dendrites. *Nature, 367*, 69–72.