To appear in: *Proceedings of the 1999 Computational Neuroscience Meeting*.
Published as part of the journal *Neurocomputing*.

Encoding Multiple Orientations in a Recurrent Network

Richard S. Zemel
Department of Psychology
University of Arizona
Tucson, AZ 85721
zemel@u.arizona.edu

Jonathan Pillow
Center for Neural Science
New York University
New York, NY
pillow@cns.nyu.edu

Abstract

Models containing recurrent connections amongst the cells within a population can account for a range of empirical data on orientation selectivity in striate cortex. However, existing recurrent models are unable to veridically encode more than one orientation at a time. Underlying this inability is an inherent limitation in the variety of activity profiles that can be stably maintained. We propose a new recurrent model that can form a broader range of stable population activity patterns. We demonstrate that these patterns preserve information about multiple orientations present in the population inputs. This preservation has significant computational consequences when information encoded in several populations must be integrated to perform behavioral tasks, such as visual discrimination.

1 Introduction

An important debate in computational neuroscience centers on the origin of selectivities in populations of cortical cells. A focus of this debate has been an extensive set of empirical data on orientation selectivity in primary visual cortex. A central question concerns how the observed sharp tuning of striate cells arises from the broadly tuned purely excitatory LGN input that grows monotonically with contrast. One class of network models posits that the sharp tuning is primarily due to the effects of recurrent connections within the striate population. These *recurrent* models can account for a wide range of the empirical data (reviewed in Sompolinsky and Shapley, 1997).

Carandini and Ringach (1997) studied a recurrent model—a simplified version of an earlier model proposed by Somers et al. (1995)—that replicates many observed characteristics of a hypercolumn of striate cells. The core of their model is a single equation implementing a feedback filter with a center-surround weighting function. Their model successfully and succinctly captures a range of results, including sharp orientation tuning and contrast-invariant tuning width. However, when the input to the cell population contains two orientations, their model exhibits peculiar responses: if the two orientations differ by less than 45°, the model responds as if the input contained a single orientation at the mean of the two values; if the angular difference between the two orientations exceeds 45°, the model responds as if they were nearly orthogonal. The model also cannot signal the presence of three orientations, and creates spurious orthogonal orientations to noisy single orientation inputs. These authors analyzed their model and showed that these effects of *attraction* and *repulsion* between orientations are unavoidable in a broad class of recurrent models of orientation selectivity.

This model thus makes strong predictions about the visual response of cortical neurons to stimuli containing multiple orientations. Relevant empirical data are scarce. DeAngelis et al. (1992) found that the response of a striate cell to its optimal orientation is reduced when the stimulus contains a second orientation. Cell responses to pairs of non-optimally oriented stimuli are not known. Nonetheless, the predictions of the model appear dubious: it predicts that the neurons will be unable to veridically encode multiple orientations, even when the orientations differ by nearly 45°. The aim of the model presented here is to explore if a broader range of activity patterns could be maintained in the population for multiple orientation stimuli. We also studied the decoding of these patterns to determine whether information about the orientations is preserved.

2 Methods

The dynamics in our model are very similar to the Carandini-Ringach (C-R) model, with two modifications. Both models contain a group of cortical neurons identical except in preferred orientation (-90°, 90°), where the response of neurons with preferred orientation ϕ to a stimulus containing orientation θ depends only on ($\theta - \phi$). Their model is governed by the following equation:

$$\tau \delta V / \delta t + V = V^{LGN} + V^{EXC} - V^{INH} \tag{1}$$

where $V(\theta, t)$ is the membrane potential of cells in orientation column θ at time t, τ is the membrane time constant, V^{LGN} is the excitatory input from LGN, and V^{EXC} and V^{INH} represent intra-cortical excitation and inhibition. These latter two quantities are obtained by convolving the responses of the cells with a narrowly tuned set of excitatory, and a broader profile of inhibitory weights, respectively. The combination of these two sets of weights produces a center-surround profile. Finally, each cell's response R is a rectified linear function of its potential V.

Our model utilizes:

1. optimized recurrent weights, both excitatory and inhibitory. The optimization procedure uses a set of examples, each consisting of the LGN input appropriate to a stimulus containing 1, 2, or 3 bars (or gratings) at random orientations, and target output values for the model striate cells. The weights are modified such that the targets are the stable activity profiles of the differential state update equation (Equation 1). The target outputs are obtained by convolving each cell's tuning function, i.e., its expected response to a single-bar stimulus as a function of its orientation, with the set of orientations present in the stimulus. These targets are based on the assumption that a cell's response to a stimulus containing a pair of oriented bars is a scaled sum of its response to the individual bars; this assumption applies to MT cell responses to multiple-direction stimuli (Recanzone et al., 1997; Treue et al., 1999).

2. nonlinear activation function. We used the function proposed by Zhang (1996): $R = a * \log(1 + \exp(b * (V + c)))^d$ (where $a, b, c, \text{and } d$ are constants) because it is easy to analytically invert, which facilitates the weight optimization procedure.

A third important addition to our model involves a different way of interpreting the activities of the striate cells. In order to determine what information about orientation is present in the model cell responses, we apply a statistical decoding method that can naturally handle multiple-valued inputs (Zemel et al., 1998). This method takes into account the cells' tuning functions in finding the set of inputs most likely to have generated the observed responses. A key feature of this method is that unimodal response profiles do not necessarily indicate that only one value is present in the input; instead, a broad unimodal profile is consistent with a stimulus containing two nearby orientations.

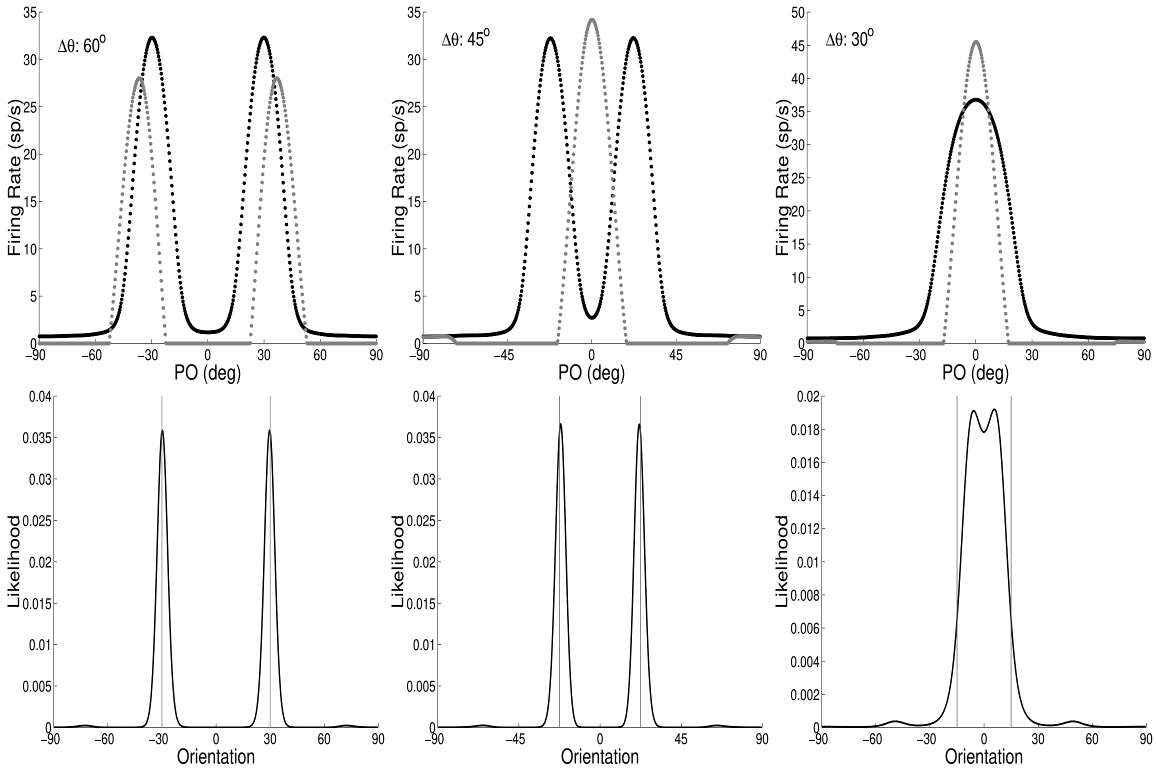

Figure 1: The top row of plots depicts stable response profiles of model cells to stimuli containing two orientations. The three columns correspond to differing angular spread between the pair of orientations. The darker dots in each plot show the cell responses in our model, the fainter dots show the responses of the C-R model to the same inputs. The bottom row of plots depicts the information about orientation decoded from those response profiles. Here the thin lines are the true orientations in the model inputs, and the thicker lines are the orientations decoded from the model responses. The plots in the central column show that our model can stably maintain bimodal activity patterns even when the modes are not orthogonal, and that information about both orientations is retained even as the angular difference shrinks below 45°.

3 Results

The response of the model to stimuli containing two orientations, of varying angular difference, is shown in Figure 1. Our model forms different stable activity patterns to all three pairs, and the decoding method accurately recovers the original orientations present in the stimulus. In contrast, the C-R model is unable to veridically encode any of these stimuli.

Our model is also able to maintain a stable response profile for stimuli containing three orientations, as well as a clean unimodal response to noisy single-orientation stimuli (see Figure 2). In both cases, the decoding method reconstructs the orientations in the stimulus.

The primary difference between the models is highlighted in Figure 3, which represents their responses to all possible two-orientation stimuli. Each horizontal slice represents the response to a particular stimulus, and the angular spread between the two orientations increases from 0°to 90°along the y-axis. The C-R model contains a discontinuity near 45°, as the population response jumps from uni- to bimodal, while our model contains a smooth transition from one to two orientations.

The optimized recurrent weights in our model (see Figure 4) have a similar shape to the center-surround weights used by most recurrent models. The wiggles in the surround of the weights gives them some power

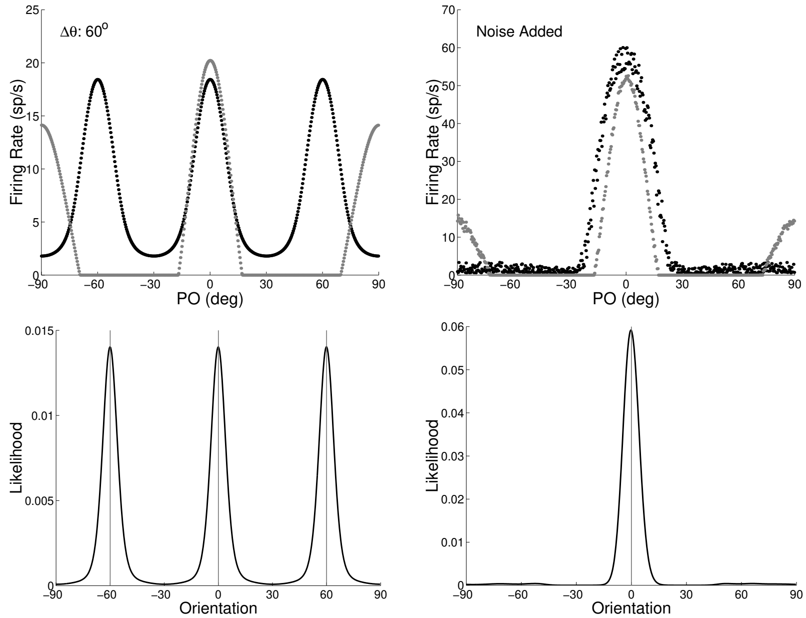

Figure 2: The response of both models to inputs containing three orientations (lines labeled as in Figure 1), and a single orientation plus noise is shown above, and the decoding of our model below.

at higher frequencies. We are currently analyzing these differences in greater detail. A crucial difference not apparent in this figure is that the relative strength of the recurrent versus feedforward weights is smaller in our model than the C-R model. This difference likely underlies our model's reduced ability to remove noise and remain strongly contrast invariant, but enhanced ability to encode multiple orientations.

4 Discussion

We have proposed a novel model of orientation selectivity in striate cortex that can support a variety of response profiles, and veridically encodes multiple orientations. Our model can replicate the cross-orientation effects observed in one of the only empirical investigations into primary visual cortex responses to multiple orientations (DeAngelis et al., 1992).

Our model makes a number of new predictions concerning responses of these cells to stimuli containing multiple orientations: (1) a range of stable activity patterns will be produced for stimuli containing two orientations with differing angular spread; (2) inputs containing three orientations separated by 60°will produce trimodal activity patterns; and (3) noisy stimuli containing single orientations will only rarely give rise to spurious bimodal response profiles.

Each of these points differ from predictions derived from the Carandini-Ringach model. They also diverge from other models of orientation selectivity, such as that proposed by Pouget et al. (1998). This model forms stable activity patterns in the population such that simple decoding techniques, such as the population vector method, closely approximate maximum likelihood estimation, which is the optimal decoding technique. This model, however, responds to pure noise stimuli as if a single orientation is present, and cannot encode multiple orientations simultaneously.

Finally, we are currently exploring the consequences of this information preservation on processing downstream from V1. It is known that V2 cells respond to illusory contours and figure-ground information, and we hypothesize that preserving information about multiple orientations within individual V1 populations plays an important role in V2 responses.

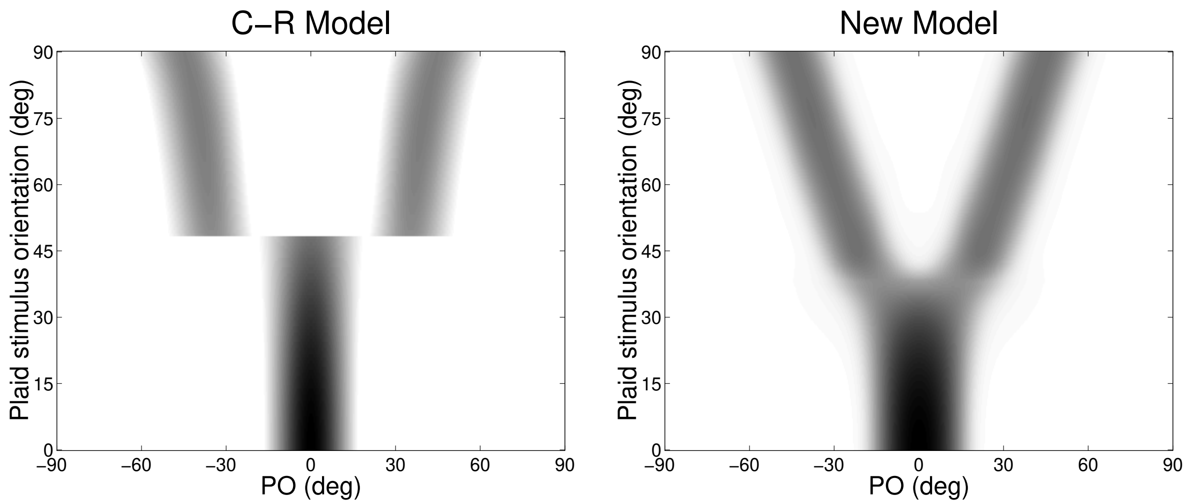

Figure 3: Summary of responses to all 2-orientation stimuli for both models. The model striate units are arranged according to their preferred orientation (PO). Each horizontal slice shows the unit activities to a particular 2-orientation stimulus. The vertical axis represents the difference between the two orientations in the stimulus, growing from 0° to 90° (maximally orthogonal) separation. Note that the C-R model response is discontinuous near 45°, while our model contains a smooth transition.

Acknowledgements: This work was funded by ONR Young Investigator Award N00014-98-1-0509 to RZ. We thank Peter Dayan and Alexandre Pouget for many useful discussions of statistical population coding.

References

[1] Carandini, M. & Ringach, D. L. (1997). Predictions of a recurrent model of orientation selectivity. Submitted to *Vision Research*.

[2] DeAngelis, G. C., Robson, J. G., Ohzawa, I., & Freeman, R. D. (1992). Organization of suppression in receptive fields of neurons in cat visual cortex. *Journal of Neurophysiology*, *68*(1), 144-163.

[3] Pouget, A., Zhang, K., Deneve, S., & Latham, P.E. (1998). Statistically efficient estimation using population codes. *Neural Computation*, *10*(2), 373-401.

[4] Somers, D. C., Nelson, S. B., & Douglas, R. J. (1995). An emergent model of orientation selectivity in cat visual cortical simple cells. *Journal of Neuroscience*, *15*, 6700-6719.

[5] Sompolinsky, H. & Shapley, R. (1997). New perspectives on the mechanisms for orientation selectivity. *Current Opinion in Neurobiology*, *7*, 514-522.

[6] Treue, S., Hol, K., & Rauber, H-J. (1999). Seeing multiple directions of motion: Physiology and psychophysics. Submitted to Nature Neuroscience.

[7] Recanzone, G. H. and Wurtz, R. H. and Schwarz, U. (1997) Responses of MT and MST neurons to one and two moving objects in the receptive field Journal of Neurophysiology, 78:6, 2904-2915.

[8] Zemel, R. S., Dayan, P., & Pouget, A. (1998). Probabilistic interpretation of population codes. *Neural Computation*, *10*, 403–430.

[9] Zhang, K. (1996). Representation of spatial orientation by the intrinsic dynamics of the head-direction cell ensemble: A theory. *Journal of Neuroscience*, *16*, 2112-2126.

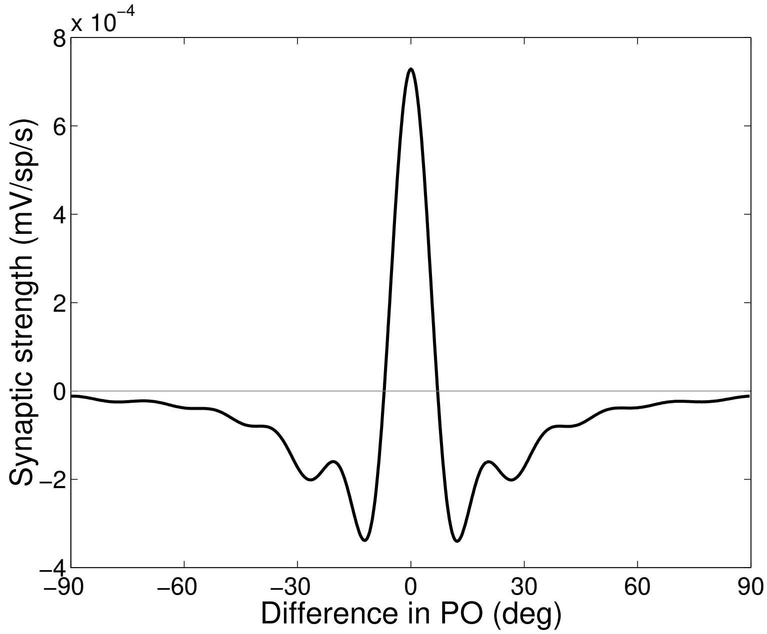

Figure 4: The optimized recurrent connection strengths between pairs of model cells in our model are plotted as a function of the difference in their preferred orientations (POs). Recurrent input is computed by convolving this weight profile with the current firing rates. The weights' center-surround profile sharpens the weak orientation bias present in LGN input without destroying information about multiple orientations.