Inferring Direction of Figure using a Recurrent Integrate-and-Fire Neural Circuit

Kyungim Baek, David H. Kim, and Paul Sajda
Department of Biomedical Engineering
Columbia University, New York, NY 10027, USA
{kb2107,dhk2102,ps629}@columbia.edu

Abstract

Several theories of early visual perception hypothesize neural circuits that are responsible for assigning ownership of an object's occluding contour to a region which represents the "figure". Previously, we presented a Bayesian network model which integrates multiple cues and uses belief propagation to infer direction of figure (DOF) along an object's occluding contour [2]. In this paper, we use a linear integrate-and-fire model to demonstrate how such inference mechanisms could be carried out in a biologically realistic neural circuit. The circuit, modeled after the network proposed by Rao [9], maps the membrane potentials of individual neurons to log probabilities and uses recurrent connections to represent transition probabilities. The network's "perception" of DOF is demonstrated for several examples, including perceptually ambiguous figures, with results qualitatively consistent with human perception.

1. Introduction

Perception requires that we infer the hidden properties (or states) of the visual scene – i.e. interpret object motion and discriminate figure-ground. The sensory input from the world, however, is inherently noisy, incomplete and ambiguous. Hidden scene properties are unlikely to be specified uniquely by a single state value, rather they are more likely to be represented by a set of possible states together with a "certainty" of each state. This naturally leads to a probabilistic representation and the machinery of probabilistic inference. In fact, it has been suggested that populations of neurons may represent and estimate probability distributions [1, 8, 14] and many researchers have formulated various aspects of human visual processing within the context of a probabilistic or Bayesian framework [4, 6, 7, 15]. In addition, numerous examples showing how Bayesian modeling can account for a variety of motion illusions [13], inference of figure-ground [12] and integration of form and motion cues for motion segmentation [11] provide strong evidence that visual integration processes are naturally defined within a Bayesian framework.

In previous work, we described a Bayesian network model for integrating spatial and/or motion cues to infer intermediate-level visual representations [2]. The model demonstrates how a locally connected network with nodes restricted to local observations can integrate visual cues for inferring scene properties. The focus of this model was on network level computation and isomorphisms with cortical circuits (e.g. columnar architecture, lateral connectivity, population responses, etc). However in this previous model the network nodes do not represent individual neurons and thus it is difficult to relate their responses and implied representations to those of real cortical neurons.

Recently, Rao has described a recurrent integrate-and-fire cortical network which can carry out Bayesian inference [9]. He suggests that neural activities may represent log-posterior probabilities and that recurrent connectivity captures the transition probabilities between states. In this paper we use this framework to implement a columnar-based neural circuit composed of integrate-and-fire neurons capable of inferring the *direction-of-figure* (DOF) in static two-dimensional scenes.

2. Inferring Direction of Figure

2.1 Network Architecture

DOF is a local representation of the ownership of an object's occluding boundary by the region which represents the object's surface [10]. It has been suggested that contour ownership is assigned at early stages of visual processing [10] and, in fact, it has been reported that more than half of the neurons in visual areas V2 and V4 are selective to contour ownership [16]. In our present computational framework, we consider the assignment of DOF as a probabilistic inference problem, with DOF being a hidden state variable that is not directly observed but can be inferred from local observations and information propagation via lateral connections.

Fig. 1 illustrates the architecture of the recurrent network for inferring DOF. The computational unit in the network simulates a cortical hypercolumn. Hypercolumns are

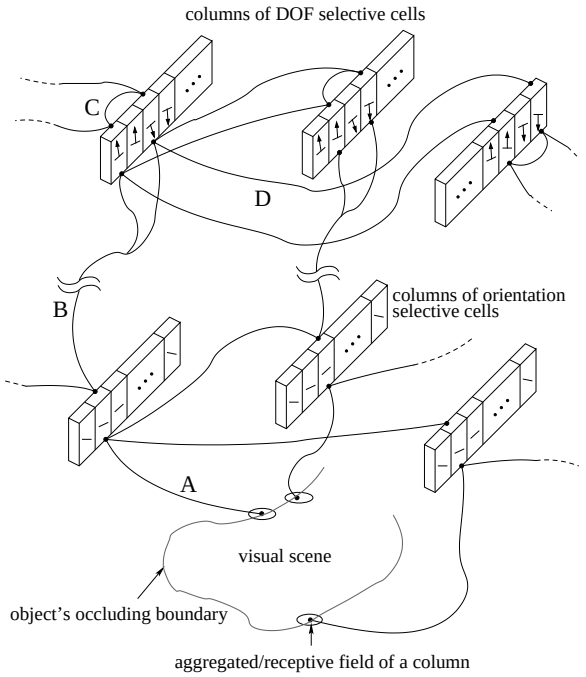

Figure 1: Architecture of the recurrent network model for inferring DOF. Information related to contour orientation and curvature is processed (via connections A, representing inputs from aggregate fields) by columns of orientation selective neurons resulting in spatial cue information that serves as the feed-forward input to DOF selective cells (via connections B). The network propagates local DOF estimates through recurrent connections (D) and integrates them until convergence. Competition exists between sets of DOF cells in the same column (via connections C).

organized laminarly and although the majority of neural connections are within columns through the layers, there exist lateral, long-range connections between sets of columns, giving rise to complex, modulatory neuronal responses [5]. In Fig. 1, the recurrent connections represent the lateral connections between columns in the superficial layers.

In the network, convexity and proximity/similarity cues are computed based on the local contour orientation passed through feed-forward and lateral connections (see [2] for details). This forms the feed-forward input to the neurons selective to DOF. There are two possible directions of figure at each location that are defined relative to the local orientation (tangent) of the contour. As can be seen in Fig. 1, the DOF selective neurons form two separate chains that locally compete with one another. The DOF is inferred by propagating and integrating local estimates based on the firing rate dynamics of the individual integrate-and-fire neurons. The firing rate of the DOF units can be interpreted as the certainty of the contour ownership in the given direction.

2.2 Bayesian Inference in a Recurrent Integrate-and-Fire Neural Circuit

We implement the columnar architecture in Fig. 1 using integrate-and-fire neurons. Let $\mathbf{u}$ be the vector of feed-forward input firing rates with synaptic weights specified by a matrix $\mathbf{W}$, and $\mathbf{v}$ be the vector of output firing rates of the recurrently connected neurons with recurrent synaptic weights specified by a matrix $\mathbf{M}$. Then the firing rate dynamics in a linear recurrent network model can be described as $\tau\frac{d\mathbf{v}}{dt} = -\mathbf{v} + \mathbf{W}\mathbf{u} + \mathbf{M}\mathbf{v}$ [3]. By discretizing this equation, v_i^t, the output firing rate of the i-th component of $\mathbf{v}$ at time t can be computed as follows:

$$v_i^t = \epsilon \mathbf{w}_i \mathbf{u}^{t-1} + \sum_j m_{ij} v_j^{t-1} \tag{1}$$

where ϵ is an integration constant, $\mathbf{w}_i$ is the i-th row vector of $\mathbf{W}$, and $m_{ij} = \epsilon M_{ij}$ for $i \neq j$ and $m_{ii} = 1 + \epsilon(M_{ij} - 1)$ [9].

Let us now consider Bayesian inference in a hidden Markov model. By applying Bayes' rule and marginalization, the posterior probability of a hidden state x_i at time t given a set of previous observations $\mathbf{y}^1, \ldots, \mathbf{y}^t$ can be computed recursively:

$$\begin{aligned} & P(x_i^t|\mathbf{y}^t, \ldots, \mathbf{y}^1) \\ &= kP(\mathbf{y}^t|x_i^t)P(x_i^t|\mathbf{y}^{t-1}, \ldots, \mathbf{y}^1) \\ &= kP(\mathbf{y}^t|x_i^t)\sum_j P(x_i^t|x_j^{t-1})P(x_j^{t-1}|\mathbf{y}^{t-1}, \ldots, \mathbf{y}^1) \end{aligned}$$

Then the posterior probability in the log domain is defined as follows:

$$\begin{aligned} \log P(x_i^t|\mathbf{y}^t, \ldots, \mathbf{y}^1) = \log k + \log P(\mathbf{y}^t|x_i^t) + \\ \log\left\{\sum_j P(x_i^t|x_j^{t-1})P(x_j^{t-1}|\mathbf{y}^{t-1}, \ldots, \mathbf{y}^1)\right\} \end{aligned} \tag{2}$$

Comparing equation (1) and equation (2), it has been shown that the recurrent integrate-and-fire network computation can implement Bayesian inference [9]. The output firing rate, weighted feed-forward input, and the recurrent term in equation (1) correspond to the log-posterior, log-likelihood, and the log of prior multiplied by the transition probabilities in equation (2) respectively:

$$v_i^t = \log P(x_i^t|\mathbf{y}^t, \ldots, \mathbf{y}^1)$$

$$\epsilon \mathbf{w}_i \mathbf{u}^{t-1} = \log P(\mathbf{y}^t|x_i^t)$$

$$\sum_j m_{ij} v_j^{t-1} = \log\left\{\sum_j P(x_i^t|x_j^{t-1})P(x_j^{t-1}|\mathbf{y}^{t-1}, \ldots, \mathbf{y}^1)\right\}$$

The recurrent weights are computed by approximating the log of sums with the sum of logs using a randomly selected

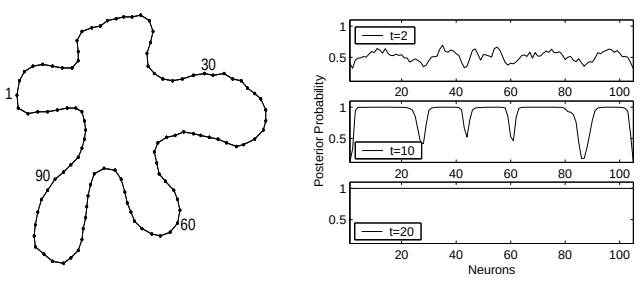

Figure 2: Posterior probabilities estimated by the DOF columnar network, shown for t=2,10 and 20 iterations, given the occluding contour shown on the left. Posterior probabilities are shown only for the neurons encoding the "correct" DOF. Indicated on the contour are indices indicating the neurons representing the DOF at the given spatial location – these indices map to those in the figures on the right. Initially, the posterior probabilities are near 0.5, indicating uncertainty in the DOF. After 20 iterations the network converges, with high certainty for all neurons encoding the correct DOF direction.

probability distribution z_j^t, which plays the role of a prior at time t:

$$\sum_j m_{ij} \log z_j^t \approx \log \left\{ \sum_j P(x_i^t | x_j^{t-1}) z_j^t \right\} \qquad (3)$$

The normalization constant $\log k$ is implemented so as to represent global recurrent inhibition.

3. Simulation Results

We apply the recurrent integrate-and-fire columnar model to several static two-dimensional figures, demonstrating its ability to infer DOF. For all the simulations presented in this section, the feed-forward weights and transition probabilities are Gaussians centered at a given spatial location. Recurrent weights are estimated from the Gaussian transition probabilities so as to minimize the squared error in equation (3), as described in [9]. Also, instead of plotting neuronal firing rates (rectified versions of the log posterior), we plot the corresponding posterior probabilities (which are more or less directly proportional to the firing rates).

Fig. 2 (left) shows an occluding contour for an arbitrarily shaped object, along with sampling points on the contour, representing the spatial locations of the columns in the network. Shown on the right are the posterior probabilities in the network for three points in time (t=2,10 and 20). Only those neurons which encode the "correct" DOF are shown (i.e. the posterior probability of the competing DOF neuron at each location is not shown). The initial DOF estimates indicate rather low certainty, as indicated by posterior

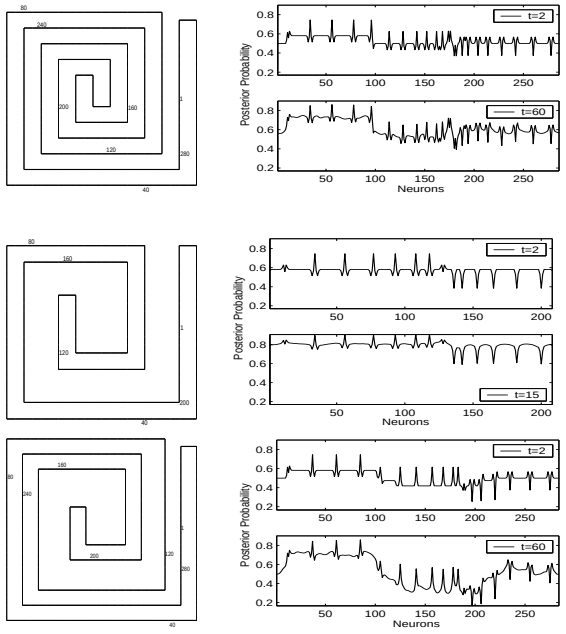

Figure 3: Integrate-and-fire columnar circuit captures the ambiguity in figure-ground perception via the posterior probabilities in the DOF. *Left*: spiral figures in which (top) figure-ground discrimination is ambiguous unless serially tracing the contour, (middle) correct discrimination of figure can be made immediately, and (bottom) increasing spiral width as it winds toward the center generates an incorrect perception of figure. *Right*: network estimates of the posterior probabilities for neurons encoding the "correct" DOF. Neurons are numbered based on indices shown on the spirals (only a subset of indices shown). For the top spiral, the posterior probabilities for the center of the spiral are approximately 0.5, indicating ambiguity in the DOF. In contrast, the low posterior probabilities for the central part of bottom spiral indicate a strong "incorrect" perception in this area of the figure. The certainty for the correct DOF in the middle spiral is very high for all locations along the contour, and in addition, the convergence rate is faster for this stimulus compared to the other two spirals (t=15 vs t=60).

probabilities around 0.5, since only local information (i.e. contour convexity) is available. However after several iterations (t=20) and propagation of activity via recurrent connections, the network quickly converges to the correct DOF, with the resulting posterior probabilities reflecting high certainty in all locations. The valleys in the middle plot correspond to the concave areas of the object contour. Because of the strong concave forward inputs, it takes more time to converge to the correct direction around those areas.

There are several classic examples in which discriminating figure from ground is not immediately clear. The top and bottom spiral figures shown in Fig. 3 are such exam-

ples [10]. Discrimination of the figure in the top spiral is difficult if we are not allowed to serially trace the contour, with difficulty increasing for regions close to the center of the spiral. On the other hand, discriminating figure in the bottom spiral seems to be straightforward. Immediately, we tend to perceive the thin strip in the center as figure, however this is in fact incorrect. In this case, the width of the spiral increases as it winds around toward the center, generating an incorrect perception of the figure-ground. Unlike these two spirals, correct figure-ground discrimination can be correctly made almost instantly for the middle spiral.

The plots on the right in Fig. 3 depict the network's output of posterior probabilities for neurons encoding the "correct" direction. The neurons are indexed as the boundary is wound toward the center and then back toward the edge, therefore the middle region of the abscissa of the plots corresponds to the central region of the spiral stimulus. The network estimates posterior probabilities at approximately 0.5 for the central part of the top spiral, indicating ambiguity of figure-ground. The posteriors are below 0.5 near the center of the bottom spiral, indicating the network "perceives" the DOF to be in the opposite, and "incorrect", direction.. For the middle spiral, the posterior probabilities are above 0.5 at all locations and mostly significantly higher than 0.5. These results illustrate the increasing ambiguity and incorrect interpretation for the central region of the top and bottom spirals, and near perfect figure-ground discrimination for the middle spiral. Furthermore, the network takes longer (i.e. twice as many iterations) to converge for the top and bottom spirals than for the middle spiral.

4. Conclusions

In this paper, we have presented a recurrent integrate-and-fire neural circuit for inferring DOF from static spatial cues. Inference in the network is based on the dynamics of neural activities, with neuronal firing rates of neurons in one of the simulated areas encoding a local representation of the DOF. Recurrent connectivity enables spatial integration of these local "beliefs". We show through simulations how the network infers DOF, including examples of perceptually ambiguous figures. Simulation results show that the network is able to account for perceptual ambiguity in DOF, qualitatively consistent with human perception.

Acknowledgments

This work was supported by the DoD Multidisciplinary University Research Initiative (MURI) program administered by the Office of Naval Research under Grant N00014-01-1-0625, and a grant from the National Imagery and Mapping Agency, NMA201-02-C0012.

References

[1] C. H. Anderson and D. C. Van Essen, "Neurobiological computational systems," in J. M. Zurada, R. J. Marks II, and C. J. Robinson (Eds.), *Computational Intelligence Imitating Life*, pp. 213-222, New York, IEEE Press, 1994.

[2] K. Baek and P. Sajda, "A probabilistic network model for integrating visual cues and inferring intermediate-level representations," In *Proceedings of IEEE Workshop on Statistical and Computational Theories of Vision*, Nice, France, October, 2003.

[3] P. Dayan and L. F. Abbott, *Theoretical Neuroscience: Computational and Mathematical Modeling of Neural Systems*, Cambridge, MA: MIT Press, 2001.

[4] W. S. Geisler and D. Kersten, "Illusions, perception and Bayes," *Nature Neuroscience*, Vol. 5, No. 6, pp. 508-510, 2002.

[5] C. D. Gilbert, "Horizontal integration and cortical dynamics," *Neuron*, Vol. 9, pp. 1-13, 1992.

[6] D. Kersten and A. Yuille, "Bayesian models of object perception," *Current Opinion in Neurobiology*, Vol. 13, pp. 1-9, 2003.

[7] T. S. Lee and D. Mumford, "Hierarchical Bayesian inference in the visual cortex," *Journal of the Optical Society of America (A)*, Vol. 20, No. 7, pp. 1434-1448, 2003.

[8] A. Pouget, P. Dayan, and R. Zemel, "Information processing with population codes," *Nature Reviews Neuroscience*, Vol. 1, pp. 125-132, 2000.

[9] R. P. N. Rao, "Bayesian computation in recurrent neural circuits," *Neural Computation*, Vol. 16, pp. 1-38, 2004.

[10] P. Sajda and L. H. Finkel, "Intermediate-level visual representations and the construction of surface perception," *Journal of Cognitive Neuroscience*, Vol. 7, No. 2, pp. 267-291, 1995.

[11] Y. Weiss and E.H. Adelson, "Perceptually organized EM: A framework for motion segmentation that combines information about form and motion," *MIT Media Lab. Perceptual Computing Section TR*, No.315, 1994.

[12] Y. Weiss, "Interpreting images by propagating Bayesian beliefs," in M. C. Mozer, M. I. Jordan, and T. Petsche (Eds.), *Advances in Neural Information Processing Systems*, Vol. 9, pp. 908-915, 1997.

[13] Y. Weiss, E.P. Simmoncelli and E.H. Adelson, "Motion illusions as optimal percepts," *Nature Neuroscience*, Vol 5, pp 598-604, 2002.

[14] R. Zemel, P. Dayan, and A. Pouget, "Probabilistic interpretation of population codes," *Neural Computation*, Vol. 10, pp. 403-430, 1998.

[15] R. Zemel, "Cortical belief networks," in R. Hecht-Neilsen, and T. McKenna (Eds.), *Computational Models for Neuroscience*, Springer-Verlag, New York, *in press*.

[16] H. Zhou, H. S. Friedman, and R. von der Heydt, "Coding of border ownership in monkey visual cortex," *Journal of Neuroscience*, Vol. 20, No. 17, pp. 6594-6611, 2000.