

Inferring Figure-Ground Using a Recurrent Integrate-and-Fire Neural Circuit

Kyungim Baek, *Member, IEEE*, and Paul Sajda, *Member, IEEE*

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 have presented a Bayesian network model which integrates multiple cues and uses belief propagation to infer local figure-ground relationships along an object’s occluding contour. 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 maps the membrane potentials of individual neurons to log probabilities and uses recurrent connections to represent transition probabilities. The network’s “perception” of figure-ground is demonstrated for several examples, including perceptually ambiguous figures, and compared qualitatively and quantitatively with human psychophysics.

Index Terms—Cortical hypercolumn, figure-ground, integrate-and-fire, probabilistic inference, visual perception.

I. INTRODUCTION

MUCH of our perceptual experience is constructed by the visual cortex. Architecturally, the visual cortex appears to be designed for integrating multiple sources of information that come through top-down, bottom-up, and horizontal connections, yet the computational mechanisms by which the integration and resulting inference occurs remain mostly unexplained. A grand challenge in neural engineering is to reverse engineer the visual cortex to reveal the essence of the computation within this distributed neural network.

Since Helmholtz described perception as being an automatic unconscious reasoning about the visual scene, researchers from various fields of study have been developing theoretical models to describe the computational mechanisms underlying visual perception. A focus has been on using probabilistic frameworks for understanding the neural mechanisms and computational principles underlying inference within the brain. Such a framework is attractive since perception can be seen as requiring inference of the hidden properties (or states) of the visual scene—e.g., interpreting object motion and discriminating figure-ground based on inherently noisy, incomplete and ambiguous sensory input. 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.

Over the past several years, it has been suggested that populations of neurons may represent and estimate probability distributions [1], [9], [16] and many researchers have formulated various aspects of human visual processing within the context of a probabilistic or Bayesian framework [4], [6], [7], [17]. In addition, numerous examples showing how Bayesian modeling can account for a variety of motion illusions [15], inference of figure-ground [14] and integration of form and motion cues for motion segmentation [13] and perception [12] provide strong evidence that visual integration processes are naturally defined within a Bayesian framework.

Though attractive as a mathematical framework, an open question is whether neurons in the brain in fact employ mechanisms for Bayesian inference. To address this question, one must consider the specifics of the anatomy and physiology in visual cortex. In previous work, we described a Bayesian network model for integrating spatial cues to infer intermediate-level representations of visual form [2] and exploit these form representations to infer scene motion [12]. The model demonstrates that 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 isomorphism with cortical circuits (Fig. 1). 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 [10]. 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 figure-ground relationship through a local representation called direction-of-figure (DOF) [11] in static two-dimensional (2-D) scenes.

II. INFERRING FIGURE-GROUND

A. Network Architecture

A visual surface can be defined by assigning ownership of an object’s occluding boundary to a region representing the object’s surface, which, therefore, determines the figure-ground

Manuscript received December 13, 2004; accepted January 13, 2005. This work was supported by the U.S. Department of Defense Multidisciplinary University Research Initiative (MURI) program administered by the Office of Naval Research under Grant N00014-01-1-0625, and by the National Imagery and Mapping Agency under Grant from NMA201-02-C0012.

The authors are with the Department of Biomedical Engineering, Columbia University, New York, NY 10027 USA (e-mail: kb2107@columbia.edu; ps629@columbia.edu).

Digital Object Identifier 10.1109/TNSRE.2005.847388

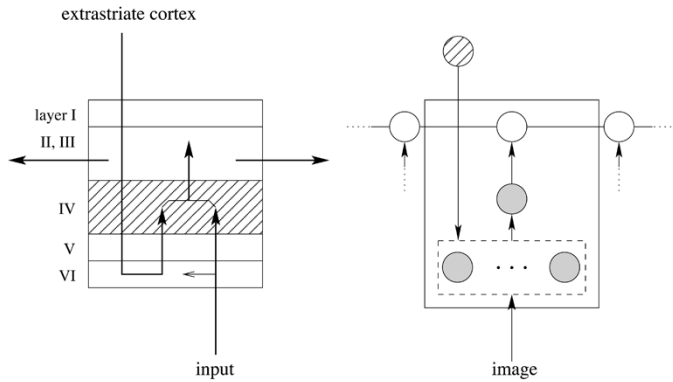


Fig. 1. Simplified diagram of cortical connections focusing on the laminar structure of a hypercolumn (left) and the isomorphic architecture used in the generative Bayesian network model (right) in our previous work. In hypercolumns, bottom-up input is provided by direct connections to layer IV and indirect connections passing through layer VI. The top-down, feedback signals from higher level cortex pass into layer VI then project into layer IV. Feed-forward connections from layer IV project to layer II/III, which forms intrinsic lateral, long-range connections between hypercolumns. In our network model, cues corresponding to bottom-up input are represented as random variables in a set of nodes (shaded circles inside the box with dashed line). Those cues are combined, together with top-down prior knowledge (circle filled with slashes), to form distributions for the observation variables (shaded circle in the middle of solid box). Observations are passed to the nodes representing hidden variables (open circles), corresponding to layers II/III in the hypercolumn structure. Spatial integration is performed by passing probabilities between neighboring hidden nodes.

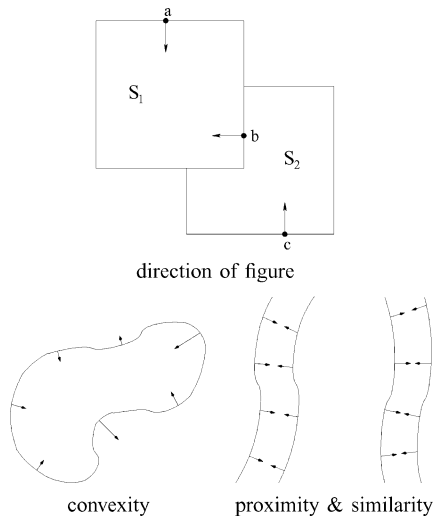


Fig. 2. DOF representation and spatial cues used in the network model. The top figure illustrates the DOF representation of the ownership of an object's surface. At location (a) the contour forms the occluding boundary of surface S_1 , i.e., surface S_1 owns the boundary. This relationship is represented by a DOF vector pointing toward surface S_1 . The local DOF is also shown for locations (b) and (c). In the network model, observed cues for the DOF include local figure convexity and proximity/similarity (bottom figures). The local convexity is determined by the local angle of the contour at a given location, while proximity/similarity cues are estimated by considering points having similar local orientation that lie in a direction orthogonal to the contour. As indicated by the length of arrows in the bottom figures, the strength of the cue is determined by the degree of convexity, similarity of local contour orientation, and the spatial distance to the similar points.

relationships [8]. DOF is a local representation of the ownership of an object's boundary [11] (Fig. 2). It has been suggested that contour ownership is assigned at early stages of visual processing [11] 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 [18]. In our present computational framework,




Fig. 3. Architecture of the recurrent network model for inferring DOF. Information related to local contour orientation and curvature in a visual image is processed by orientation selective neurons, resulting in spatial cue information that serves as the feed-forward input to DOF selective neurons. Since the visual world is retinotopically mapped onto the cortical surface (i.e., neighboring neurons are activated by the stimuli in adjacent positions in the visual image) and a neuron sees only a limited local area of visual field, communication between neighboring neurons carries out spatial propagation and integration of the local information. The network propagates local DOF estimates through recurrent connections (A) and integrates them until convergence. Competition exists between sets of DOF neurons in the same column (B).

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 propagated via lateral connections.

Fig. 3 illustrates the architecture of the recurrent neural circuit for inferring DOF. The computational unit in the network simulates a DOF selective neuron in a cortical hypercolumn, encoding DOF relative to the local orientation (tangent). Hypercolumns are organized laminarly and although the majority of neural connections are within columns, there exist lateral, long-range connections between sets of columns, giving rise to complex, modulatory neuronal responses [5]. In Fig. 3, the recurrent connections represent the lateral connections between columns of DOF selective neurons 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 (Fig. 2; 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 of the contour. As can be seen in Fig. 3, 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.

B. Bayesian Inference in a Recurrent Integrate-and-Fire Neural Circuit

We implement the columnar architecture in Fig. 3 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 $\tau(d\mathbf{v}/dt) = -\mathbf{v} + \mathbf{W}\mathbf{u} + \mathbf{M}\mathbf{v}$ [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

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

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 $m_{ii} = 1 + \epsilon(M_{ii} - 1)$ [10].

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 $\mathbf{y}^1, \dots, \mathbf{y}^t$ can be computed recursively as

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

where k is the normalization constant. Then, the posterior probability in the log domain is defined as

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

Comparing (1) and (2), it has been shown that the recurrent integrate-and-fire network computation can implement Bayesian inference [10]. The output firing rate, weighted feed-forward input, and the recurrent term in (1) correspond to the log-posterior, log-likelihood, and the log of prior multiplied by the transition probabilities in (2) respectively, namely

$$\begin{aligned} v_i^t &= \log P(x_i^t | \mathbf{y}^t, \dots, \mathbf{y}^1) \\ \epsilon \mathbf{w}_i \mathbf{u}^{t-1} &= \log P(\mathbf{y}^t | x_i^t) \\ \sum_j m_{ij} v_j^{t-1} &= \log \left\{ \sum_j P(x_i^t | x_j^{t-1}) P(x_j^{t-1} | \mathbf{y}^{t-1}, \dots, \mathbf{y}^1) \right\}. \end{aligned}$$

The recurrent weights are computed by approximating the log of sums with the sum of logs using a randomly selected probability distribution z_j^t , which plays the role of a prior at time t

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

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

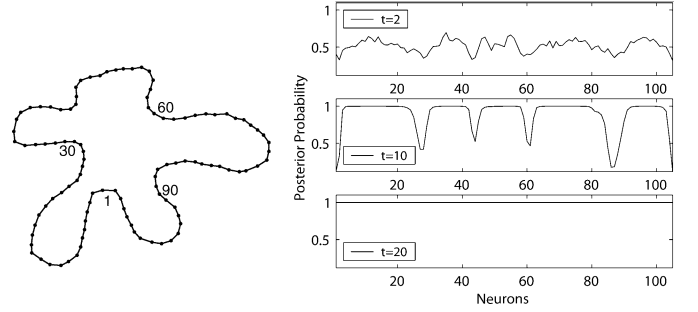


Fig. 4. Posterior probabilities estimated by the DOF columnar network, shown for $t = 2, 10$, and 20 iterations, given the occluding contour shown on the left. Posterior probabilities are shown only for neurons encoding the correct DOF. On the contour are indexes indicating the neurons representing the DOF at the given spatial location—these indexes 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.

III. SIMULATION RESULTS

We apply the recurrent integrate-and-fire columnar model to several static 2-D 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 (3), as described in [10]. 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. 4 (left) shows an occluding contour for an arbitrarily shaped object, along with sampling points on the contour, representing the corresponding 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 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. Those locations with strong concavity cues, communicated via feed-forward connections, have a slower convergence to the correct figure-ground direction.

There are several classic examples in which discriminating figure-ground is ambiguous. The top and bottom spiral figures shown in Fig. 5 are such examples [11]. 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, the discriminating figure in the bottom spiral seems more 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




Fig. 5. 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 indexes shown on the spirals (only a subset of indexes 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 ($t = 15$ versus $t = 60$).

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. 5 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 net-




Fig. 6. The 2-AFC experiment for psychophysically determining a subject's performance for correctly discriminating figure-ground direction. In each trial, a stimulus containing one of the three spirals shown in Fig. 5 is presented for 0.5 s and the subject is instructed to report the figure-ground direction relative to a specified point on the contour of the spiral. As soon as the subject makes a response a new fixation cross appears for 1 s at which time the next trial begins. The locations on the contour in which subjects are asked to discriminate figure-ground are shown in Fig. 7. In any given trial the spiral and location along the contour are chosen at random (three spirals and eight locations). In addition, the orientation of the spiral is randomized in order to reduce learning effects (one of four orientations).

work "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., four times as many iterations) to converge for the top and bottom spirals, compared to the middle spiral.

IV. COMPARISON TO THE PSYCHOPHYSICAL RESULTS

To understand how well the network's "perception" of DOF matches human perception, we designed a psychophysical experiment for discriminating figure-ground direction using a two alternative forced choice (2-AFC) paradigm. In each trial, a stimulus consisting of a spiral and a dot located on the spiral's contour is presented for 0.5 s after a 1-s fixation period. The task is to determine the figure-ground direction relative to the dot location. In all cases the figure-ground direction is on the left or right side of the dot. The subject is instructed to report their decision as quickly as possible by pressing either the left or right mouse button. The screen remains blank after the 0.5-s presentation of a stimulus and then, after the subject's response, a new fixation cross appears for 1 s, indicating the start of the next trial (Fig. 6). For each spiral, a dot can appear at one of eight randomly selected locations, as shown on the left in Fig. 7. In addition, the orientation of each spiral is randomized across four different configurations—original, left-right flip, and 180° rotation of the previous two. Thus, the stimulus set consists of a total of 96 images from which a stimulus is drawn randomly in each trial.

Six naive subjects who were not familiar with the stimuli participated in the experiments. Results were obtained from nine sessions of the experiment since three of the subjects participated for the second time about ten months after their first participation. Each session starts with nine practice trials followed by four test blocks containing 96 trials in each block. Since a stimulus is randomly drawn without replacement, each image in the stimulus set appears exactly once in a test block. In each experiment, the subject’s discrimination performance for figure-ground direction at a given dot location on a spiral is estimated from the 16 responses [four configurations and four presentations (blocks)]. We consider this the psychophysical correlate of the DOF.

The network performance is computed by averaging posterior probabilities for the correct DOF at the three nearest neighbors of the corresponding dot location used in the psychophysical experiments. The three plots on the right in Fig. 7 show nine percentage correct values of human subjects at each dot location of the three spirals along with the model performances represented by asterisks. For the bottom spiral the network’s perception is within the range specified by the spread of subjects’ perception at all dot locations, while for half of the cases it falls outside of the range for the top and middle spirals.

For each dot location, we performed a t-test to determine whether the network performance was significantly different from the mean subject performance (indicated by squares in Fig. 7). The results show that the network’s estimate of the DOF is not significantly different from the mean subject performance for one (location 3), three (location 4, 5, and 8) and zero cases for the top, middle and bottom spirals respectively. Thus, while in 15/24 cases the network performance falls within the spread of the subjects’ performance, only in 4/24 cases is it not significantly different from the mean subject’s performance.

There are several factors that may explain the quantitative difference between the model’s performance and human psychophysics. First, using only convexity and proximity/similarity cues for inferring DOF may limit the network’s performance. For example, use of “closure,” one of the perceptual cues for DOF, would bias the network’s perception more toward the incorrect figure direction in the central region (location 2, 3, 6, and 7) of the middle spiral and more toward the correct figure direction at location 1 in all spirals, thereby reducing the difference in performance between the human subjects and the model.

Another factor which may account for the difference between the model and human subjects’ performance is related to subjects’ adaptation to the task and strategy/learning during the experiment. Many subjects report that they develop a strategy for discriminating figure direction near the edge of the spiral boundary—i.e., starting with the “figure,” they alternate regions, labeling every other one as figure. The influence of such a strategy can be seen for location 1 for all three spirals and locations 4, 5, and 8 for the top spiral where the mean subject performance is much higher than the network performance. Use of such a strategy also explains the high variance of subjects’ responses at locations 4 and 8 for the middle spiral where the figure direction tends to be biased toward the narrow strip. This bias appears to be reduced for those subjects who employ such a high-level strategy. Learning and employing such high-level

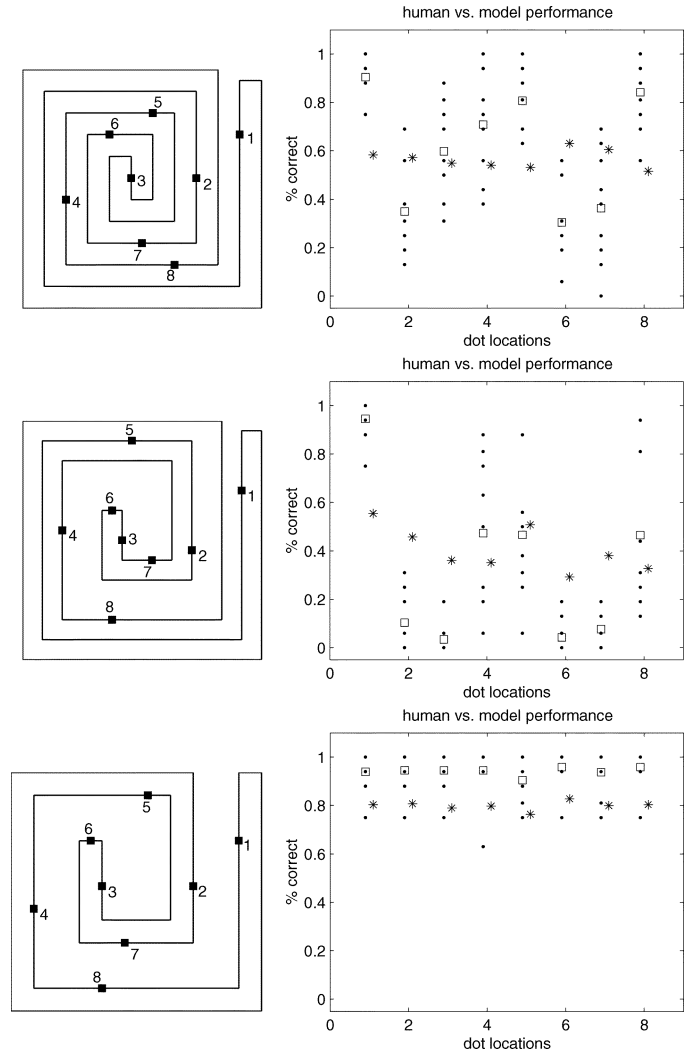


Fig. 7. Eight dot locations on the contour of the spirals used in the psychophysical experiments. The plots on the right show nine percentage correct values computed from 16 responses at each dot location in a session of psychophysical experiments (dots) along with the posterior probabilities for correct DOF of the network model (asterisks). The mean subject performance across the nine sessions is represented by squares. Note that the percent correct values for a dot location can be less than nine because of the overlap between the values from different sessions. The network’s posterior probabilities shown in the plots are computed by averaging those at the three nearest neighbors of the corresponding dot locations.

strategies is beyond the scope of the current model and can be seen as a top-down input into our network, potentially simulated as a prior cue in the model [2].

V. CONCLUSION

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 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. Further comparison

with the performance of human subjects leads to several factors that potentially contribute to the quantitative difference between the model's performance and psychophysics.

Although the model we present is simple and far from complete, it provides insight into whether the human visual system possesses the necessary neural machinery for inferring figure-ground and locally representing the results of this inference, including the underlying uncertainty. However, there are several questions for further investigation. One that is particularly intriguing is what biologically plausible learning rules could be used to estimate the recurrent and feed-forward synaptic weights directly from input data. While the network model we propose is based on the dynamics of neural activities of cortical neurons, the synaptic weights in the network are set as simple Gaussians. Learning these synaptic weights from the input data would be more consistent with the idea that the visual system exploits the statistics of natural images to infer properties about the scene.

ACKNOWLEDGMENT

The authors would like to thank D. Kim for help in implementing and testing the model.

REFERENCES

- [1] C. H. Anderson and D. C. Van Essen, "Neurobiological computational systems," in *Computational Intelligence Imitating Life*, J. M. Zurada, R. J. Marks II, and C. J. Robinson, Eds. New York: IEEE Press, 1994, pp. 213–222.
- [2] K. Baek and P. Sajda, "A probabilistic network model for integrating visual cues and inferring intermediate-level representations," in *Proc. IEEE Workshop Statistical and Computational Theories of Vision*, Nice, France, Oct. 2003, http://www.stat.ucla.edu/~yuille/meetings/2003_workshop.php.
- [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 Neurosci.*, 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 Neurobiol.*, vol. 13, pp. 1–9, 2003.
- [7] T. S. Lee and D. Mumford, "Hierarchical Bayesian inference in the visual cortex," *J. Opt. Soc. Amer., A*, vol. 20, no. 7, pp. 1434–1448, 2003.
- [8] K. Nakayama, S. Shimojo, and G. H. Silverman, "Stereoscopic depth: Its relation to image segmentation, grouping, and the recognition of occluded objects," *Perception*, vol. 18, pp. 55–68, 1989.
- [9] A. Pouget, P. Dayan, and R. Zemel, "Information processing with population codes," *Nature Rev. Neurosci.*, vol. 1, pp. 125–132, 2000.
- [10] R. P. N. Rao, "Bayesian computation in recurrent neural circuits," *Neural Computation*, vol. 16, pp. 1–38, 2004.
- [11] P. Sajda and L. H. Finkel, "Intermediate-level visual representations and the construction of surface perception," *J. Cogn. Neurosci.*, vol. 7, no. 2, pp. 267–291, 1995.
- [12] P. Sajda and K. Baek, "Integration of form and motion within a generative model of visual cortex," *Neural Netw.*, vol. 17, pp. 809–821, 2004.
- [13] 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.
- [14] Y. Weiss, "Interpreting images by propagating Bayesian beliefs," in *Advances in Neural Information Processing Systems*, M. C. Mozer, M. I. Jordan, and T. Petsche, Eds. Cambridge, MA: MIT Press, 1997, vol. 9, pp. 908–915.
- [15] Y. Weiss, E. P. Simoncelli, and E. H. Adelson, "Motion illusions as optimal percepts," *Nature Neurosci.*, vol. 5, pp. 598–604, 2002.
- [16] R. Zemel, P. Dayan, and A. Pouget, "Probabilistic interpretation of population codes," *Neural Computation*, vol. 10, pp. 403–430, 1998.
- [17] R. Zemel, "Cortical belief networks," in *Computational Models for Neuroscience*, R. Hecht-Neilsen and T. McKenna, Eds. New York: Springer-Verlag, 2003, pp. 267–287.
- [18] H. Zhou, H. S. Friedman, and R. von der Heydt, "Coding of border ownership in monkey visual cortex," *J. Neurosci.*, vol. 20, no. 17, pp. 6594–6611, 2000.



Kyungim Baek (M'05) received the B.S. degree (with highest distinction) in mathematics and the M.S. degree in computer science from Sogang University, Seoul, Korea, in 1994 and 1996, and the Ph.D. degree in computer science from Colorado State University, Fort Collins, in 2002.

Presently, she is a Postdoctoral Research Fellow in the Laboratory for Intelligent Imaging and Neural Computing (LIINC), the Department of Biomedical Engineering, Columbia University, New York. In 1996, she was a Research Fellow with the Computer

Graphics Laboratory, Sogang University. Her research interests include computational models of visual perception, Bayesian methods, machine learning, neural computation, and related applications in computer vision.

Dr. Baek received the Best Cognitive Vision Paper Award at the International Conference on Vision Systems 2003.



Paul Sajda (S'91–M'93) received the B.S. degree in electrical engineering from the Massachusetts Institute of Technology (MIT), Cambridge, in 1989, and the M.S. and Ph.D. degrees in bioengineering from the University of Pennsylvania (UPenn), Philadelphia, in 1992 and 1994, respectively.

In 1994, he joined the David Sarnoff Research Center, Princeton, NJ, where he went on to become the Head of the Adaptive Image and Signal Processing Group. He is currently an Associate Professor of biomedical engineering and radiology at Columbia University, New York, NY, where he is Director of the Laboratory for Intelligent Imaging and Neural Computing (LIINC). Currently, he has over 80 publications and holds 5 U.S. patents. His research interests include neural engineering, neuroimaging, computational neural modeling and machine learning applied to image understanding.

Dr. Sajda is an Associate Editor for the IEEE TRANSACTIONS ON BIOMEDICAL ENGINEERING. He has received several awards for his research including: the Adler Award for outstanding undergraduate research in Electrical Engineering (MIT, 1989); the Flexner Award for outstanding thesis research in the Neurosciences (UPenn, 1993); the Pollack Award for outstanding dissertation research in Bioengineering (UPenn, 1994); the Sarnoff Technical Achievement Award (1996); and a National Science Foundation CAREER Award (2002).