Bayesian Computation in Recurrent Cortical Circuits

Rajesh P. N. Rao
Department of Computer Science and Engineering
University of Washington
Seattle, WA 98195

E-mail: *rao@cs.washington.edu*

Technical Report No. 02-07-04

CSE Department, University of Washington

July, 2002

Abstract

A large number of human psychophysical results have been successfully explained in recent years using Bayesian models. However, the neural implementation of such models remains largely unclear. In this paper, we show that a network architecture commonly used to model the cerebral cortex can implement Bayesian inference for an arbitrary Markov model. We illustrate the suggested approach using a visual motion detection task. Our simulation results show that the model network exhibits direction selectivity and correctly computes the posterior probabilities for motion direction. When used to solve the well-known random dots motion discrimination task, the model generates responses that mimic the activities of evidence-accumulating neurons in cortical areas LIP and FEF. In addition, the model predicts reaction time distributions that are similar to those obtained in human psychophysical experiments that manipulate the prior probabilities of targets and task urgency.

1 Introduction

Bayesian models of perception have proved extremely useful in explaining results from a range of human psychophysical experiments [Knill and Richards, 1996, Rao et al., 2002]. These psychophysical tasks range from inferring 3D structure from 2D images and judging depth from multiple cues to perception of motion and color (see [Bloj et al., 1999, Weiss et al., 2002] and chapters by Mamassian *et al*. and Jacobs in [Rao et al., 2002]). The strength of the Bayesian approach lies in its ability to quantitatively model the interaction between prior knowledge and sensory evidence. Bayes rule prescribes how prior probabilities and stimulus likelihoods should be combined, allowing the responses of subjects to be interpreted in terms of the resulting posterior distributions.

Additional support for Bayesian models (also known as "ideal observer" models) comes from recent neurophysiological and psychophysical studies on visual decision making. Carpenter and colleagues have shown that the reaction time distribution of human subjects making eye movements to one of two targets is well-explained by a model computing the log likelihood ratio of one target over the other [Carpenter and Williams, 1995]. In another study, the saccadic response time distribution of monkeys could be predicted from the time taken by neural activity in area FEF to reach a fixed threshold [Hanes and Schall, 1996], suggesting that these neurons are involved in accumulating evidence (interpreted as log likelihoods) over time. Similar activity has also been reported in the primate area LIP [Shadlen and Newsome, 2001]. A mathematical model based on log likelihood ratios was found to be consistent with the observed neurophysiological results [Gold and Shadlen, 2001].

Given the growing evidence for Bayesian models in perception and decision

making, an important question is how such models could be implemented within the recurrent networks of the primate brain. Previous models of probabilistic inference in cortical networks have been based on the concept of neural tuning curves or kernel functions and have focused on the estimation of static quantities such as stimulus location [Anderson and Van Essen, 1994, Zemel et al., 1998, Deneve et al., 1999, Pouget et al., 2000]. Other models have relied on mean-field approximations or various forms of Gibbs sampling for perceptual inference [Hinton and Sejnowski, 1986, Dayan et al., 1995, Dayan and Hinton, 1996, Hinton and Ghahramani, 1997, Rao and Ballard, 1997, Rao, 1999, Rao and Ballard, 1999, Hinton and Brown, 2002].

We describe a new approach to Bayesian computation in a cortical network model. We specify how the feedforward and recurrent connections in the network may be selected to perform Bayesian inference for an arbitrary hidden Markov model. The approach is illustrated using a visual motion detection task. The direction selective model neurons are shown to correctly compute the log posterior probability of stimulus location and direction. When used in a random dots motion discrimination task, the model produces responses that are qualitatively similar to the responses of "evidence accumulating" neurons in primate brain areas LIP and FEF.

2 Modeling Cortical Networks

We begin by considering a commonly used architecture for modeling cortical response properties, namely, a recurrent network with firing-rate dynamics (see, for example, [Dayan and Abbott, 2001]). Let $\mathbf{u}$ denote the vector of input firing rates to the network and let $\mathbf{v}$ represent the output firing rates of the recurrently connected neurons in the network. Let $\mathbf{W}$ represent the feedforward synaptic weight matrix and $\mathbf{M}$ the recurrent

weight matrix. The following equation describes the dynamics of the network:

$$\tau \frac{d\mathbf{v}}{dt} = -\mathbf{v} + \mathbf{W}\mathbf{u} + \mathbf{M}\mathbf{v} \tag{1}$$

The equation can be rewritten in a discrete form as follows:

$$v_i(t+1) = v_i(t) + \epsilon(-v_i(t) + \mathbf{w}_i\mathbf{u}(t) + \sum_j m_{ij}v_j(t)) \tag{2}$$

$$\text{i.e.} \quad v_i(t+1) = \epsilon\mathbf{w}_i\mathbf{u}(t) + \sum_j m'_{ij}v_j(t) \tag{3}$$

where ϵ is the integration rate, $m'_{ii} = 1 + \epsilon(m_{ii} - 1)$ and for $i \neq j$, $m'_{ij} = \epsilon m_{ij}$. How can Bayesian computation be achieved using the dynamics given by Equation 3? To answer this question, we consider Bayesian inference in a simple Markov model.

3 Bayesian Computation in a Cortical Network

Let $\theta(t)$ represent the hidden state of a Markov model at time t with transition probabilities given by $P(\theta(t) = \theta_i|\theta(t-1) = \theta_j)$ for $i, j = 1 \ldots N$. Let $\mathbf{I}(t)$ be the observable output governed by the probabilities $P(\mathbf{I}(t)|\theta(t))$. Then, the posterior probability of being in state θ_i at time t is given by Bayes rule:

$$P(\theta(t) = \theta_i|\mathbf{I}(t), \ldots, \mathbf{I}(1)) = kP(\mathbf{I}(t)|\theta(t) = \theta_i)P(\theta(t) = \theta_i|\mathbf{I}(t-1), \ldots, \mathbf{I}(1)) \tag{4}$$

where k is a normalization constant. The equation can be made recursive with respect to the posterior probabilities as follows:

$$P(\theta_i^t|\mathbf{I}(t), \ldots, \mathbf{I}(1)) = kP(\mathbf{I}(t)|\theta_i^t) \sum_j P(\theta_i^t|\theta_j^{t-1})P(\theta_j^{t-1}|\mathbf{I}(t-1), \ldots, \mathbf{I}(1)) \tag{5}$$

where we have written "$\theta(t) = \theta_i$" as θ_i^t for notational convenience. This equation can be written in the log domain as:

$$\log P(\theta_i^t|\mathbf{I}(t), \ldots, \mathbf{I}(1)) = \log P(\mathbf{I}(t)|\theta_i^t) + \log k + \log\Big[\sum_j P(\theta_i^t|\theta_j^{t-1})P(\theta_j^{t-1}|\mathbf{I}(t-1), \ldots, \mathbf{I}(1))\Big] \tag{6}$$

A recurrent network governed by Equation 3 can implement Equation 6 if:

$$v_i(t+1) = \log P(\theta_i^t|\mathbf{I}(t),\ldots,\mathbf{I}(1)) \quad (7)$$

$$\epsilon \mathbf{w}_i \mathbf{u}(t) = \mathbf{F}(\theta_i)\mathbf{I}(t) = \log P(\mathbf{I}(t)|\theta_i^t) \quad (8)$$

$$\sum_j m'_{ij} v_j(t) = \log\Big[\sum_j P(\theta_i^t|\theta_j^{t-1}) P(\theta_j^{t-1}|\mathbf{I}(t-1),\ldots,\mathbf{I}(1))\Big] \quad (9)$$

with the normalization term $\log k$ being added after each update. In the above, we have assumed that the log likelihoods are given by a linear filter $\mathbf{F}(\theta_i)$ ($= \epsilon \mathbf{w}_i$). For example, $\mathbf{F}(\theta_i)$ could represent a spatially localized Gaussian or Gabor filter when the inputs $\mathbf{I}(t)$ are images.

A more challenging problem is to pick the recurrent weights m'_{ij} such that Equation 9 holds true. This amounts to approximating a *log sum* with a *sum of logs*. We find the weights m'_{ij} by generating a set of random probabilities $x_j(t)$ for $t = 1,\ldots,M$ and solving the system of equations given by:

$$\sum_j m'_{ij} \log x_j(t) = \log\Big[\sum_j P(\theta_i^t|\theta_j^{t-1}) x_j(t)\Big] \quad (10)$$

for all i and t. A solution can be computed using the standard pseudoinverse method to produce an appropriate set of recurrent weights m'_{ij} for any given set of transition probabilities $P(\theta_i^t|\theta_j^{t-1})$ (see, for example, Figures 1 and 2).

4 An Example Task: Detecting Visual Motion

We illustrate the proposed approach using the example of visual motion detection. A prominent property of visual cortical cells is selectivity to the direction of visual motion. In this section, we interpret the activity of such cells as representing the log posterior of motion direction at a spatial location given a series of input images.

For simplicity, we focus on the case of 1-D motion in an image consisting of P pixels with two possible motion directions: leftward (L) or rightward (R). Let the state

θ_{ij} represent motion at a spatial location i in direction $j \in \{L, R\}$. Consider a network of N neurons, each representing a particular state θ_{ij} (Figure 1A). The feedforward weights are Gaussians i.e. $\mathbf{F}(\theta_{iR}) = \mathbf{F}(\theta_{iL}) = \mathbf{F}(\theta_i) =$ Gaussian centered at location i with a standard deviation σ. Figure 1B depicts the feedforward weights for a network of 30 neurons, 15 encoding leftward and 15 encoding rightward motion.

To detect motion, the transition probabilities $P(\theta_{ij}|\theta_{kl})$ must be selected to reflect both the direction of motion and speed of the moving stimulus. For the present study, the transition probabilities for rightward motion from the state θ_{kR} (i.e. $P(\theta_{iR}|\theta_{kR})$) were set according to a Gaussian centered at location $k + x$, where x is a parameter determined by stimulus speed. The transition probabilities for leftward motion from the state θ_{kL} were likewise set to Gaussian values centered at $k - x$. The transition probabilities from states near the two boundaries ($i = 1$ and $i = P$) were chosen to be uniformly random values. Figure 1C shows the matrix of transition probabilities. Given these transition probabilities, the recurrent weights m'_{ij} can be computed using Equation 10 (see Figure 1D). Figure 2 shows that the log sum in Equation 10 can indeed be approximated quite accurately by a sum of logs using the weights m'_{ij}.

5 Results

5.1 Bayesian Estimation of Motion

Figure 3 shows the output of the network in the middle of a sequence of input images depicting a bar moving either leftward or rightward. As shown in the figure, for a leftward moving bar at position i, the highest network output is for the neuron representing location i and direction L, while for a rightward moving bar, the neuron representing location i and direction R has the highest output. The outputs are plotted as the log posteriors plus a constant offset to allow interpretation as nonnegative firing rates. Note that

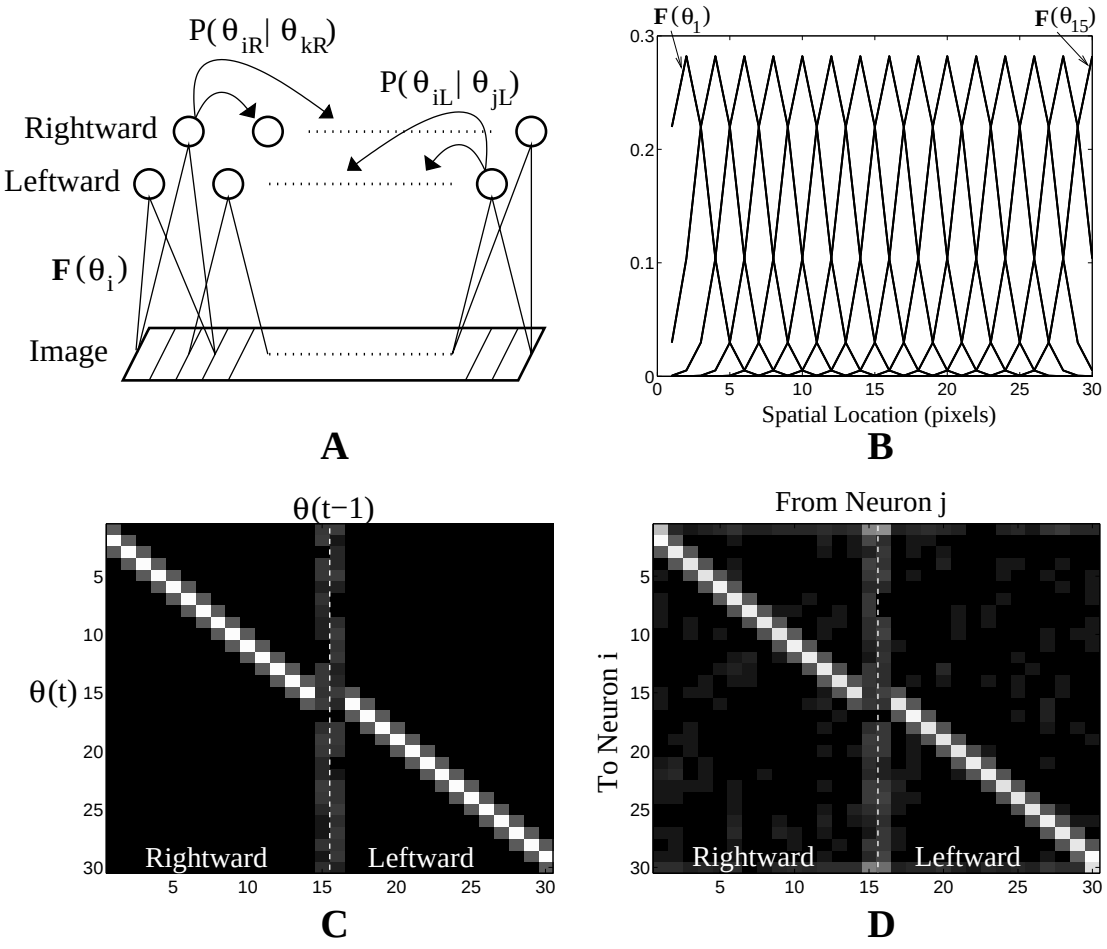

Figure 1: **Recurrent Network for Motion Detection**. (A) depicts a recurrent network of neurons, shown for clarity as two chains selective for leftward and rightward motion respectively. The feedforward synaptic weights for neuron i (in the leftward or rightward chain) are determined by $\mathbf{F}(\theta_i)$. The recurrent weights reflect the transition probabilities $P(\theta_{iR}|\theta_{kR})$ and $P(\theta_{iL}|\theta_{jL})$. (B) Feedforward weights $\mathbf{F}(\theta_i)$ for neurons $i = 1, \ldots, 15$ (rightward chain). The feedforward weights for neurons $i = 15, \ldots, 30$ (leftward chain) are identical. (C) Transition probabilities $P(\theta(t) = \theta_i|\theta(t-1) = \theta_j)$ for $i = 1, \ldots, 30$. Probability values are proportional to pixel brightness. (D) Recurrent weights m'_{ij} computed from the transition probabilities in (C) using Equation 10.

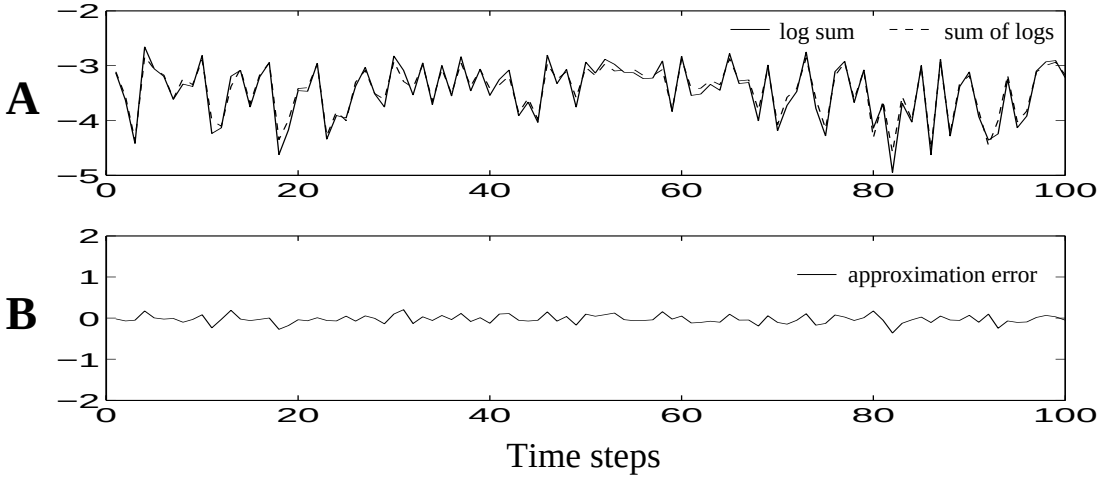

Figure 2: **Approximating Transition Probabilities with Recurrent Weights**. (A) compares the right and left sides of Equation 10 for random test data $x_j(t)$ (different from the training $x_j(t)$) using the transition probabilities and weights shown in Figures 1C and 1D respectively. The solid line represents the right side of Equation 10 while the dotted line represents the left side. Only the results for $i = 8$ are shown but similar results were obtained for the rest of the neurons. (B) shows the approximation error (difference between the two sides of Equation 10 over time).

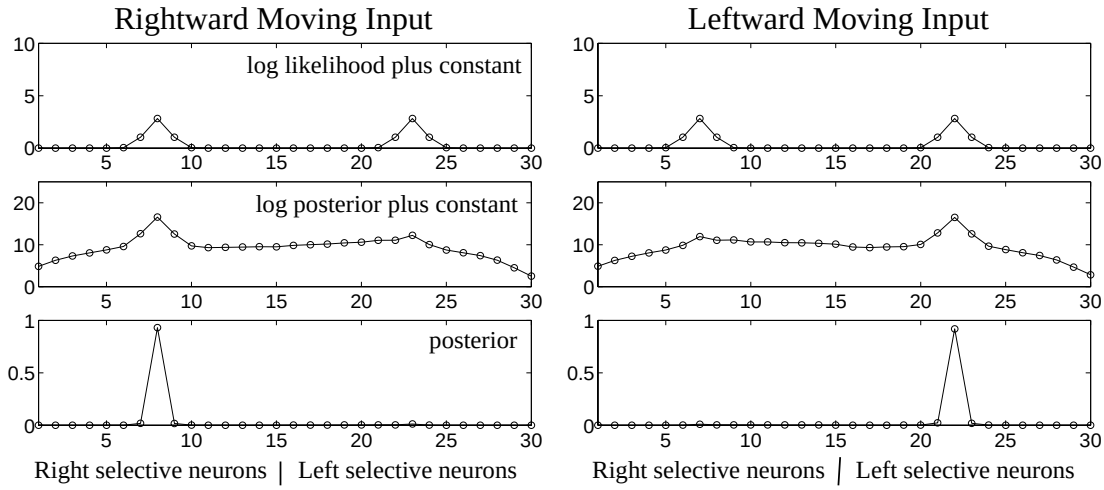

Figure 3: **Network Output for a Moving Stimulus**. (Left Panel) The three plots depict the log likelihood, log posterior, and posterior probability for a rightward moving bar when it arrives at the central image location. The outputs of the 30 neurons in the network are given by the middle graph. Constant offsets have been added to the log likelihood and log posteriors to allow interpretation as nonnegative firing rates. Note that the log likelihoods are the same for the left and right selective neurons (the first 15 and last 15 neurons respectively, as dictated by the recurrent weights in Figure 1D) but their outputs correctly reflect the direction of motion due to recurrent interactions. (Right Panel) The same three plots for a leftward moving bar as it reaches the central location.

even though the log likelihoods are the same for leftward and rightward moving inputs, the asymmetric recurrent weights (which represent the transition probabilities) allow the network to distinguish between leftward and rightward moving stimuli. The plot of posteriors in Figure 3 shows that the network correctly computes posterior probabilities close to 1 for the states θ_{iL} and θ_{iR} for leftward and rightward motion respectively at location i.

5.2 The Random Dots Task

To establish a connection to behavioral data, we consider the well-known random dots motion discrimination task (see, for example, [Shadlen and Newsome, 2001]). The stimulus consists of an image sequence showing a group of moving dots, a fixed fraction of which is randomly selected at each frame and moved in a fixed direction (for example, either left or right). The rest of the dots are moved in random directions. The fraction of dots moving in the same direction is called the *coherence* of the stimulus. Figure 4A depicts the stimulus for two different levels of coherence. The task is to decide the direction of motion of the coherently moving dots for a given input sequence. A wealth of data exists on the psychophysical performance of humans and monkeys as well as the neural responses in brain areas such as MT and LIP in monkeys performing the task (see [Shadlen and Newsome, 2001] and references therein). Our goal is to explore the extent to which the proposed model can explain the existing data for this task and make testable predictions.

The motion detection network in the previous section can be used to decide the direction of coherent motion by computing the posterior probabilities for left and right motion given the input images. These probabilities can be computed by marginalizing the posterior distribution computed by the neurons for left (L) and right (R) motion over

all spatial positions i:

$$P(L|\mathbf{I}(t),\ldots,\mathbf{I}(1)) = \sum_i P(\theta_{iL}|\mathbf{I}(t),\ldots,\mathbf{I}(1)) \quad (11)$$

$$P(R|\mathbf{I}(t),\ldots,\mathbf{I}(1)) = \sum_i P(\theta_{iR}|\mathbf{I}(t),\ldots,\mathbf{I}(1)) \quad (12)$$

Note that the log of these marginal probabilities can also be computed directly from the log posteriors (i.e. outputs of the neurons) using the log sum approximation method described in Section 3. The log posterior ratio $r(t)$ of leftward over rightward motion can be used to decide the direction of motion:

$$r(t) = \log P(L|\mathbf{I}(t),\ldots,\mathbf{I}(1)) - \log P(R|\mathbf{I}(t),\ldots,\mathbf{I}(1)) \quad (13)$$

$$= \log \frac{P(L|\mathbf{I}(t),\ldots,\mathbf{I}(1))}{P(R|\mathbf{I}(t),\ldots,\mathbf{I}(1))} \quad (14)$$

If $r(t) > 0$, the evidence seen so far favors leftward motion and vice versa for $r(t) < 0$. The instantaneous ratio $r(t)$ is susceptible to rapid fluctuations due to the noisy stimulus. We therefore use the following decision variable $d_L(t)$ to track the running average of the log posterior ratio of L over R:

$$d_L(t+1) = d_L(t) + \alpha(r(t) - d_L(t)) \quad (15)$$

and likewise for $d_R(t)$. We assume that these variables are computed by a separate set of "decision neurons" that receive inputs from the motion detection network. The output of the model is "L" if $d_L(t) > c$ and "R" if $d_R(t) > c$, where c is a threshold parameter that depends on task constraints (for example, accuracy versus speed requirements) [Reddi and Carpenter, 2000].

Figures 4B and 4C show the responses of the two decision neurons over time for two different directions of motion and two levels of coherence. Besides correctly computing the direction of coherent motion in each case, the model also responds faster when the

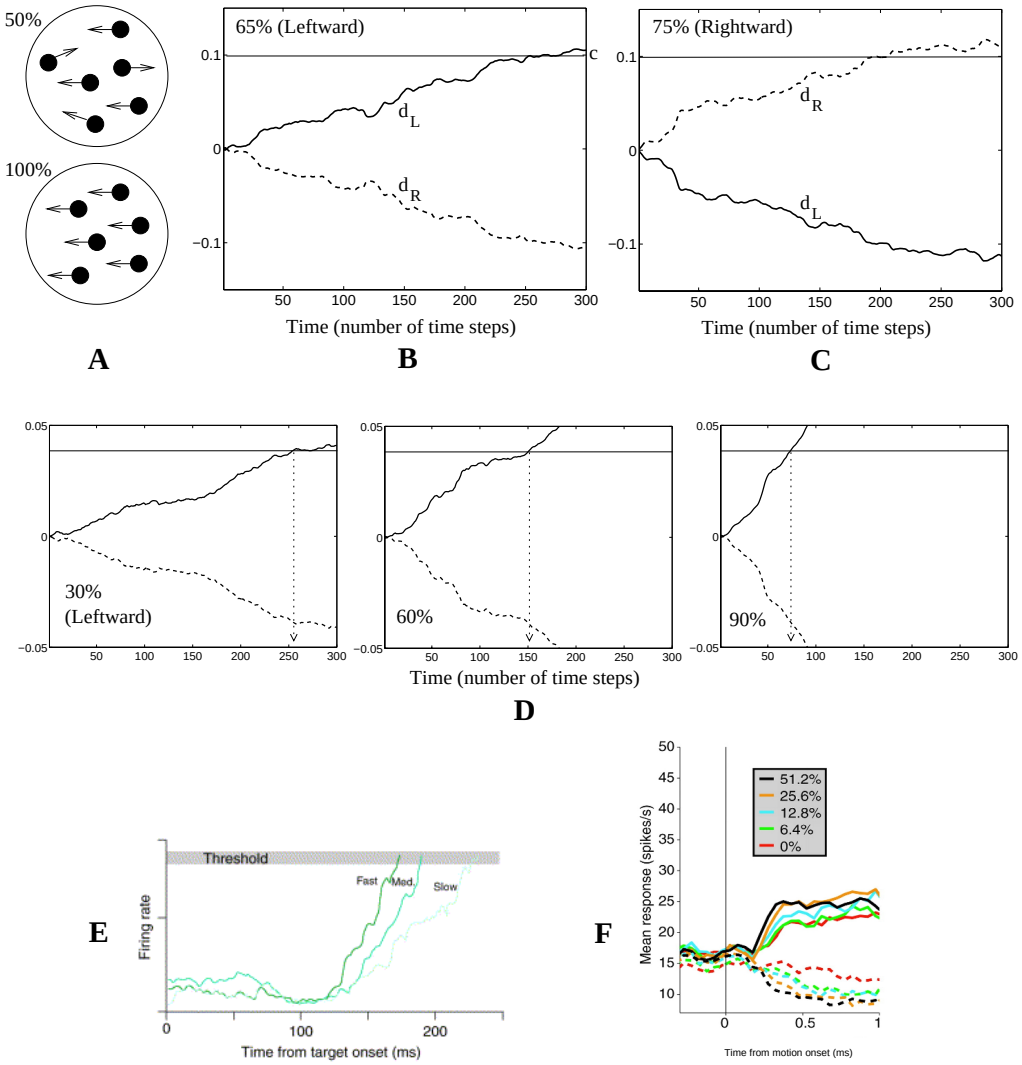

Figure 4: **Output of Decision Neurons in the Model**. (A) Depiction of the random dots task. Two different levels of motion coherence (50% and 100%) are shown. A 1-D version of this stimulus was used in the model simulations. (B) & (C) Outputs $d_L(t)$ and $d_R(t)$ of model decision neurons for two different directions of motion. The decision threshold is labeled "c". (D) Outputs of decision neurons for three different levels of motion coherence. Note the increase in rate of evidence accumulation at higher coherencies. For a fixed decision threshold, the model predicts faster reaction times for higher coherencies (dotted arrows). (E) Responses of a neuron in area FEF for a monkey performing an eye movement task (based on [Hanes and Schall, 1996, Schall and Hanes, 1998]). Faster reaction times were associated with a more rapid rise to a fixed threshold (shown as three different colored activity profiles). (F) Averaged firing rate of 104 neurons in area LIP for the random dots task (from [Shadlen and Newsome, 2001]). Solid and dashed curves represent trials in which the monkey judged motion direction toward and away from the receptive field of the neuron, respectively. The slope and magnitude of the response are both affected by motion coherence (see different colored activity profiles) in a manner that resembles the model responses shown in (D). Unlike our simulations and the FEF task in (E), the monkey in this case was trained to wait through a delay period before responding, causing the LIP responses to plateau rather than terminate after reaching a decision threshold.

stimulus has higher coherence. This phenomenon can be appreciated more clearly in Figure 4D, which predicts progressively shorter reaction times for increasingly coherent stimuli (dotted arrows).

5.3 Comparison to Neurophysiological Data

The relationship between faster rates of evidence accumulation and shorter reaction times has received experimental support from a number of studies. Figure 4E shows the responses of a neuron in the frontal eye fields (FEF) for fast, medium, and slow responses to a visual target [Schall and Hanes, 1998]. Schall and collaborators have shown that the distribution of monkey response times can be reproduced using the time taken by neural activity in FEF to reach a fixed threshold [Hanes and Schall, 1996]. A similar rise-to-threshold model by Carpenter and colleagues has received strong support in human psychophysical experiments that manipulate the prior probabilities of targets [Carpenter and Williams, 1995] and the urgency of the task [Reddi and Carpenter, 2000] (see next section).

In the case of the random dots task, Shadlen and collaborators have shown that in primates, one of the cortical areas involved in making the decision regarding coherent motion direction is area LIP. The activities of many neurons in this area progressively increase during the motion viewing period, with faster rates of rise for more coherent stimuli (see Figure 4F). This behavior is similar to the responses of "decision neurons" in the model (Figure 4B–4D), suggesting that the outputs of the recorded LIP neurons could be interpreted as representing the log posterior ratio of one task alternative over another (see [Carpenter and Williams, 1995, Gold and Shadlen, 2001] for related suggestions).

6 Model Predictions

The strength of a Bayesian model lies in its ability to predict the effects of varying prior knowledge or various types of biases in a given task. We explored the consequences of varying two types of biases in the motion discrimination task: (a) changing the prior probability of a particular motion direction (L or R) and (b) lowering the decision threshold to favor speed instead of accuracy.

Changing the prior probability for a particular motion direction amounts to initializing the decision neuron activities d_L and d_R to the appropriate log prior ratios:

$$d_L(0) = \log P(L) - \log P(R) \quad (16)$$

$$d_R(0) = \log P(R) - \log P(L) \quad (17)$$

where $P(L)$ and $P(R)$ denote the prior probabilities of left and right motion directions respectively. Note that for the equiprobable case, these variables get initialized to 0 as in the simulations presented thus far. We examined the distribution of reaction times obtained as a result of assuming a slightly higher prior probability for L compared to R (0.51 versus 0.49). Figure 5A compares this distribution to the distribution obtained for the unbiased (equiprobable) case (both distributions are restricted to trials where the direction of motion was L). The model predicts that the entire distribution for a particular task choice should shift leftwards (towards shorter reaction times) when the subject assumes a higher prior for that particular choice.

Similarly, when time is at a premium, a subject may opt to lower the decision threshold to reach decisions quickly, albeit at the risk of making some incorrect decisions. Such an urgency constraint was simulated in the model by lowering the decision threshold to half the original value, keeping the stimulus coherency constant (60% in this case). As shown in Figure 5B, the model again predicts a leftward shift of the entire

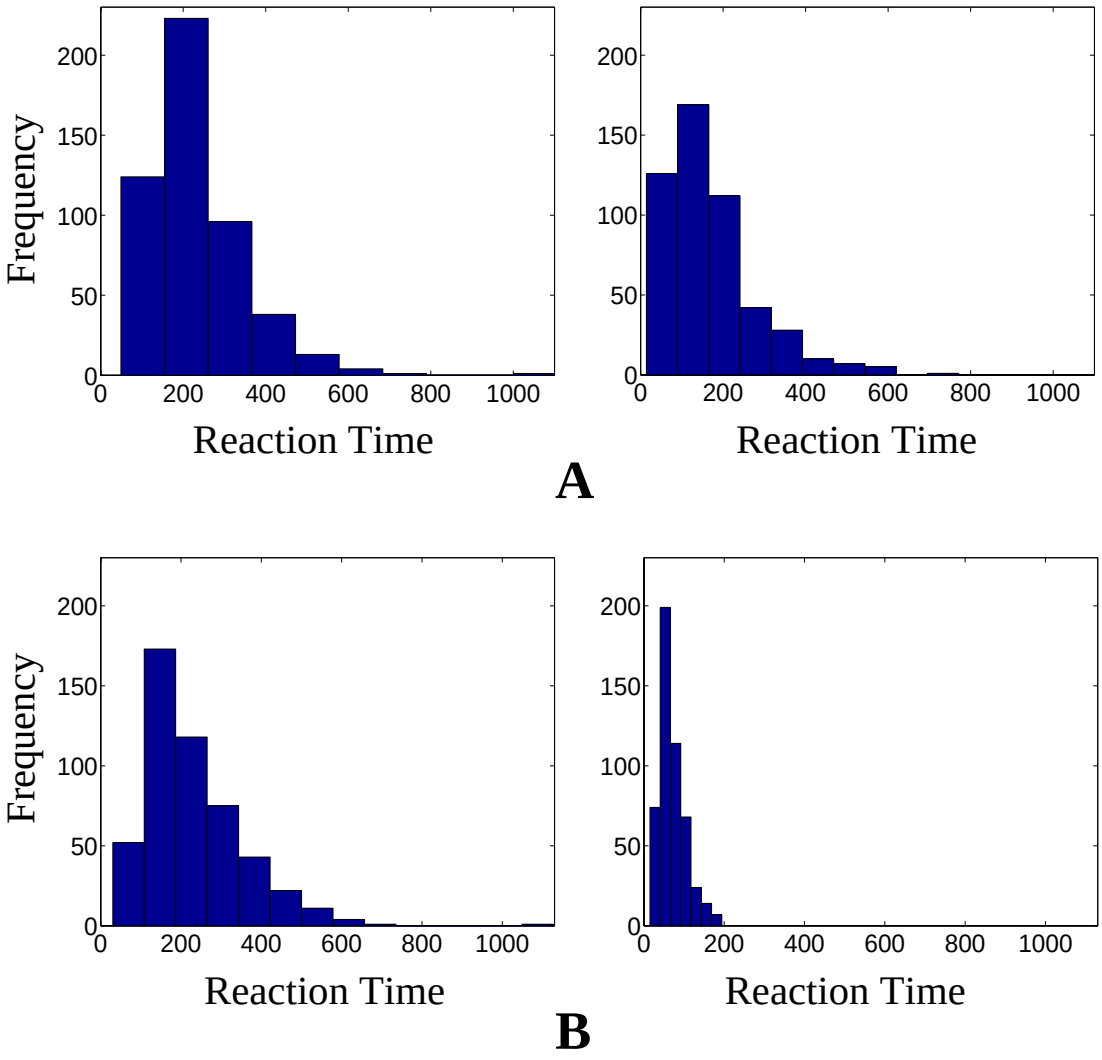

Figure 5: **Reaction Time Distributions as a Function of Prior Probability and Urgency**. (A) The left panel shows the distribution of reaction times for the motion discrimination task obtained from model simulations when leftward and rightward motion are equiprobable. The right panel shows the distribution for the case where leftward motion is more probable than rightward motion (0.51 versus 0.49). Note the leftward shift of the distribution towards shorter reaction times. Both distributions here were restricted to leftward motion trials only. (B) The right panel shows the result of halving the decision threshold, simulating an urgency constraint that favors speed over accuracy. Compared to the non-urgent case (left panel), the reaction time distribution again exhibits a leftward shift towards faster (but potentially inaccurate) responses.

reaction time histogram, corresponding to faster (but sometimes inaccurate) responses.

Although these predictions are yet be tested in the context of the motion discrimination task (see [Mazurek et al., 2001] for some preliminary results), psychophysical experiments by Carpenter and colleagues [Carpenter and Williams, 1995] lend some support to the model's predictions. In these experiments, a human subject was required to make an eye movement to a dim target that appeared randomly to the left or right of a fixation point after a random delay of 0.5–1.5 seconds. The latency distributions of the eye movements were recorded as a function of different prior probabilities for one of the targets. For convenience, these distributions were plotted as cumulative histograms on a reciprobit scale: the X-axis was 1/latency while the Y-axis used a probit scale. Thus, the skewed latency histograms map to straight lines in the reciprobit plots. Figure 6A shows how a shift in the latency histogram results in a swiveling of the reciprobit plot. Figure 6B shows the reciprobit plots of response latencies for a subject as a function of target prior probabilities. As predicted by the model for the motion discrimination task (Figure 5A), the response latency histogram shifts leftwards for increasing prior probabilities, resulting in a leftward rotation of the reciprobit plots. The faster response time distributions in the lower part of the plot have been attributed to "express saccades" which cannot be accounted for by the present model.

In a second set of experiments [Reddi and Carpenter, 2000], Reddi and Carpenter examined the effects of imposing urgency constraints on the eye movement task. They asked subjects to respond as rapidly as possible and to worry less about the accuracy of their response. As shown in Figure 6C, the reciprobit plot swivels to the left for the urgency condition, indicating that the latency histogram has shifted leftwards as suggested by the model simulation results in Figure 5B. Furthermore, in a remarkable follow-up experiment, Reddi and Carpenter showed that this leftward shift in the his-

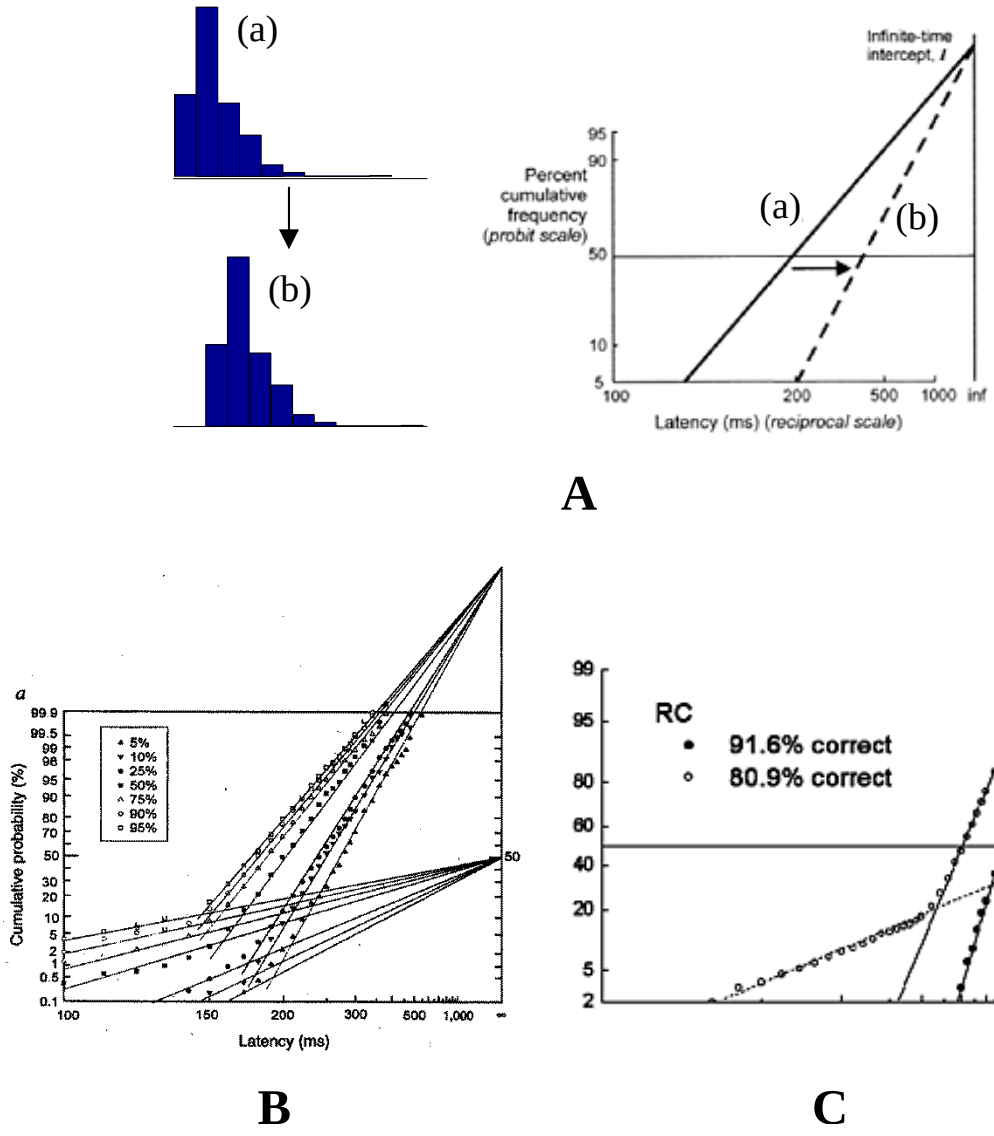

Figure 6: **Effects of Changing Prior Probability and Task Urgency in an Eye Movement Task** (based on [Carpenter and Williams, 1995, Reddi and Carpenter, 2000]). (A) illustrates the effect of a rightward shift in the reaction time distribution (from (a) to (b)) on the corresponding reciprobit plot (right panel; see text for details). (B) Response latency distributions for a subject in the eye movement task (plotted on a reciprobit scale) as a function of target prior probabilities. Note the leftward rotation for higher priors. (C) Response latency distributions for a subject in the eye movement task, when the subject was instructed to respond as fast as possible (white circles) or as accurately as possibly (dark circles). Note the leftward swiveling of the distribution for the urgency condition (corresponding to a lowering of the decision threshold in the model).

togram due to urgency could be countered by a corresponding rightward shift due to lower prior probability for one of the targets, resulting in no overall shift in the latency histogram [Reddi and Carpenter, 2000]. This provides strong evidence for a rise-to-threshold model of decision making, where the threshold is chosen according to the constraints of the task at hand.

7 Summary and Conclusions

We have shown that a recurrent network of leaky integrator neurons can approximate Bayesian inference for an arbitrary Markov model. Such networks have previously been used to model a variety of cortical response properties [Dayan and Abbott, 2001]. Our results suggest a new interpretation of neural firing rates in such networks as representing log posterior probabilities. Using the example of visual motion detection, we showed how such networks can exhibit direction selectivity and compute the log posterior probability of motion direction given a sequence of input images. We also showed how these outputs could be used to decide the direction of coherent motion in the random dots task. The activity of the "decision neurons" in the model resemble the activity of evidence-accumulating neurons in LIP and FEF, two cortical areas known to be involved in visual decision making. The model correctly predicts reaction times as a function of stimulus coherence, and produces reaction time distributions that are qualitatively similar to those obtained in eye movement experiments that manipulate prior probabilities of targets and task urgency.

Although the results obtained thus far are encouraging, several important questions remain. For example, how can the feedforward and recurrent weights be learned directly from input data? We intend to explore this question using biologically-plausible approximations to the expectation-maximization (EM) algorithm

[Dempster et al., 1977]. A related question is how the suggested framework could be extended to allow hierarchical Bayesian inference in multi-layer cortical networks. Previous models of hierarchical inference and learning in cortical circuits should prove useful in addressing this question [Dayan et al., 1995, Dayan and Hinton, 1996, Hinton and Ghahramani, 1997, Rao and Ballard, 1997, Rao and Ballard, 1999]. A final question of interest is how the rewards associated with various choices available in a task influence the decision neurons in the model. We expect ideas from reinforcement learning and Bayesian decision theory as well as recent neurophysiological results in the monkey [Platt and Glimcher, 1999] to be helpful in guiding our research in this direction.

Acknowledgments

This research was supported by a Sloan Research Fellowship and NSF grant no. 130705. I am grateful to Aaron Hertzmann, Michael Shadlen, Michael Rudd, and John Palmer for discussions and comments.

References

[Anderson and Van Essen, 1994] Anderson, C. H. and Van Essen, D. C. (1994). Neurobiological computational systems. In Zurada, J. M., II, R. J. M., and Robinson, C. J., editors, *Computational Intelligence: Imitating Life*, pages 213–222. New York, NY: IEEE Press.

[Bloj et al., 1999] Bloj, M. G., Kersten, D., and Hurlbert, A. C. (1999). Perception of three-dimensional shape influences colour perception through mutual illumination. *Nature*, 402:877–879.

[Carpenter and Williams, 1995] Carpenter, R. H. S. and Williams, M. L. L. (1995). Neural computation of log likelihood in control of saccadic eye movements. *Nature*, 377:59–62.

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

[Dayan and Hinton, 1996] Dayan, P. and Hinton, G. (1996). Varieties of Helmholtz machine. *Neural Networks*, 9(8):1385–1403.

[Dayan et al., 1995] Dayan, P., Hinton, G., Neal, R., and Zemel, R. (1995). The Helmholtz machine. *Neural Computation*, 7:889–904.

[Dempster et al., 1977] Dempster, A. P., Laird, N. M., and Rubin, D. B. (1977). Maximum likelihood from incomplete data via the EM algorithm. *J. Royal Statistical Society Series B*, 39:1–38.

[Deneve et al., 1999] Deneve, S., Latham, P. E., and Pouget, A. (1999). Reading population codes: a neural implementation of ideal observers. *Nature Neuroscience*, 2(8):740–745.

[Gold and Shadlen, 2001] Gold, J. I. and Shadlen, M. N. (2001). Neural computations that underlie decisions about sensory stimuli. *Trends in Cognitive Sciences*, 5(1):10–16.

[Hanes and Schall, 1996] Hanes, D. P. and Schall, J. D. (1996). Neural control of voluntary movement initiation. *Science*, 274:427–430.

[Hinton and Sejnowski, 1986] Hinton, G. and Sejnowski, T. (1986). Learning and relearning in Boltzmann machines. In Rumelhart, D. and McClelland, J., editors, *Parallel Distributed Processing*, volume 1, chapter 7, pages 282–317. MIT Press, Cambridge.

[Hinton and Brown, 2002] Hinton, G. E. and Brown, A. D. (2002). Learning to use spike timing in a restricted Boltzmann machine. In *Probabilistic Models of the Brain: Perception and Neural Function*, pages 285–296. Cambridge, MA: MIT Press.

[Hinton and Ghahramani, 1997] Hinton, G. E. and Ghahramani, Z. (1997). Generative models for discovering sparse distributed representations. *Phil. Trans. Roy. Soc. Lond. B.* To appear.

[Knill and Richards, 1996] Knill, D. C. and Richards, W. (1996). *Perception as Bayesian Inference*. Cambridge, UK: Cambridge University Press.

[Mazurek et al., 2001] Mazurek, M. E., Ditterich, J., Palmer, J., and Shadlen, M. N. (2001). Effect of prior probability on behavior and LIP responses in a motion discrimination task (abstract no. 58.11). *Society for Neuroscience Abstracts*, 27.

[Platt and Glimcher, 1999] Platt, M. L. and Glimcher, P. W. (1999). Neural correlates of decision variables in parietal cortex. *Nature*, 400:233–238.

[Pouget et al., 2000] Pouget, A., Dayan, P., and Zemel, R. S. (2000). Information processing with population codes. *Nature Reviews Neuroscience*, 1(2):125–132.

[Rao, 1999] Rao, R. P. N. (1999). An optimal estimation approach to visual perception and learning. *Vision Research*, 39(11):1963–1989.

[Rao and Ballard, 1997] Rao, R. P. N. and Ballard, D. H. (1997). Dynamic model of visual recognition predicts neural response properties in the visual cortex. *Neural Computation*, 9(4):721–763.

[Rao and Ballard, 1999] Rao, R. P. N. and Ballard, D. H. (1999). Predictive coding in the visual cortex: A functional interpretation of some extra-classical receptive field effects. *Nature Neuroscience*, 2(1):79–87.

[Rao et al., 2002] Rao, R. P. N., Olshausen, B. A., and Lewicki, M. S. (2002). *Probabilistic Models of the Brain: Perception and Neural Function*. Cambridge, MA: MIT Press.

[Reddi and Carpenter, 2000] Reddi, B. A. and Carpenter, R. H. (2000). The influence of urgency on decision time. *Nature Neuroscience*, 3(8):827–830.

[Schall and Hanes, 1998] Schall, J. D. and Hanes, D. P. (1998). Neural mechanisms of selection and control of visually guided eye movements. *Neural Networks*, 11:1241–1251.

[Shadlen and Newsome, 2001] Shadlen, M. N. and Newsome, W. T. (2001). Neural basis of a perceptual decision in the parietal cortex (area LIP) of the rhesus monkey. *Journal of Neurophysiology*, 86(4):1916–1936.

[Weiss et al., 2002] Weiss, Y., Simoncelli, E. P., and Adelson, E. H. (2002). Motion illusions as optimal percepts. *Nature Neuroscience*, 5(6):598–604.

[Zemel et al., 1998] Zemel, R. S., Dayan, P., and Pouget, A. (1998). Probabilistic interpretation of population codes. *Neural Computation*, 10(2):403–430.