Deciphering the Neural Code for Retinal Ganglion Cells through Statistical Inference

by

Yi-Chieh Wu

B.S.E.E., Rice University (2007)

Submitted to the Department of Electrical Engineering and Computer Science
in partial fulfillment of the requirements for the degree of

Master of Science in Electrical Engineering and Computer Science

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

June 2009

Author .... 5/22/09 ...
Department of Electrical Engineering and Computer Science
May 22, 2009

Certified by... .... 5/22/2009 ...
John L. Wyatt
Professor
Thesis Supervisor

Accepted by ... .
Terry P. Orlando
Chairman, Department Committee on Graduate Theses

Deciphering the Neural Code for Retinal Ganglion Cells through Statistical Inference

by

Yi-Chieh Wu

Submitted to the Department of Electrical Engineering and Computer Science
on May 22, 2009, in partial fulfillment of the
requirements for the degree of
Master of Science in Electrical Engineering and Computer Science

Abstract

This work studies how the visual system encodes information in the firing patterns of retinal ganglion cells. We present a visual scene to a retina, obtain in-vitro recordings from a multi-electrode array, and attempt to identify or reconstruct the scene. Our approach uses the well-known linear-nonlinear Poisson model to characterize neural firing behavior and accounts for stochastic variability by fitting parameters using maximum likelihood. To characterize cells, we use white noise analysis followed by numerical optimization to maximize the likelihood of the experimentally observed neural responses. We then validate our method by keeping these fitted parameters constant and using them to estimate the speed and direction of moving edges, and to identify a natural scene out of a set of possible candidates. Limitations of our approach, including reconstruction fidelity and the validity of various assumption are also examined through simulated cell responses.

Thesis Supervisor: John L. Wyatt
Title: Professor

Acknowledgments

To my advisor John Wyatt for his support, time, and dedication, and for providing me with the opportunity to merge my interests of engineering and biology.

To my academic advisor Asu Ozdaglar for helping me navigate the world of MIT EECS.

To my undergraduate advisors Richard Baraniuk and Lydia Kavraki for giving me some of my first opportunities for research and for motivating me to continue my education in graduate school.

To my officemates Stavros and Shamim. Stavros, for taking the brave first steps in starting up the group work in neural coding, and Shamim, for all the discussions and keeping me company during those long hours in the office.

To my collaborators Steve, Ofer, and Shelley, for their help in conducting experiments and for sharing their vast knowledge of retinal physiology.

To my friends Han, Pavitra, Bradley, and Becky for keeping me sane, whether through talking about problem sets or watching *BSG*.

Finally, to my parents Rongsong and Manhung, and my sisters, Emmy and Diana, for their love, support, and guidance.

---

[0] This material is based upon work supported under a National Science Foundation Graduate Research Fellowship.

Contents

List of Figures

List of Tables

Nomenclature

Model

LNP	linear-nonlinear Poisson
MISE	mean integrated square error
ML	maximum likelihood
STA	spike-triggered average

Neural Physiology

ISI	interspike interval
LED	local edge detector
PSTH	peristimulus time histogram
RF	receptive field
RGC	retinal ganglion cell

Miscellaneous

MEA	multi-electrode array

Chapter 1

Introduction

The neural system represents and transmits information through a complicated network of cells and interconnections so that we can perceive the world around us, and a major challenge in neuroscience is understanding how this encoding and decoding takes place. Besides increasing our knowledge of perception, this neural coding problem is also of interest to researchers for its host of potential applications. Among others, it has been hypothesized that the human body adapted to code natural stimuli in an efficient, error-resilient manner, and understanding how sensory stimuli are represented in neural responses could provide a basis for developing better coding algorithms. Another major application involves reconstructing a person's experience using neural activity to study differences in perception or understand how the senses are integrated in the perceptual pathway. Perhaps the most direct application is towards helping people interact with their environment, as seen in the success of the cochlear implant and other neural prosthetics.

In this thesis, we focus on one small part of this larger problem: *visual neural coding*. We consider the problem as one of image identification and reconstruction: given neural activity, determine the input image. We propose algorithms for a visual decoder using the responses from a population of retinal ganglion cells (RGCs), and as experimental verification, we stimulate rabbit retina with a visual stimulus, record responses using a multi-electrode array (MEA), and attempt to either identify or reconstruct the scene. We focus on a small subset of cell types known as ON and

OFF cells and analyze the encoding of artificial and natural stimuli.

RGCs are chosen in part because this work is done in collaboration with the Boston Retinal Implant Project, whose goal is to develop an implantable microelectronic prosthesis to restore vision to people with degenerative eye conditions such as retinitis pigmentosa or age-related macular degeneration (see Figure 1-1). These diseases cause the loss of photoreceptors in the outer retina while leaving the ganglion cell layer almost entirely functional[19]. The ganglion cells are the only retinal cells that feed signals to the brain, and this connection is only feed-forward[8]. Theoretically, this implies that if we are able to replicate the spiking response caused by a light pattern in the ganglion cells of a healthy retina via electrode stimulation, we could effectively create visual perceptions. Furthermore, by focusing efforts on a relatively early portion of the visual pathway, we take advantage of higher-level processing in the visual cortex and allow the implant to be minimally invasive, as it is external to the eye so that the retina remains intact. The neural coding problem is of further interest in terms of an implant as such an understanding would provide an objective metric in order to judge the response to electrical stimulation.

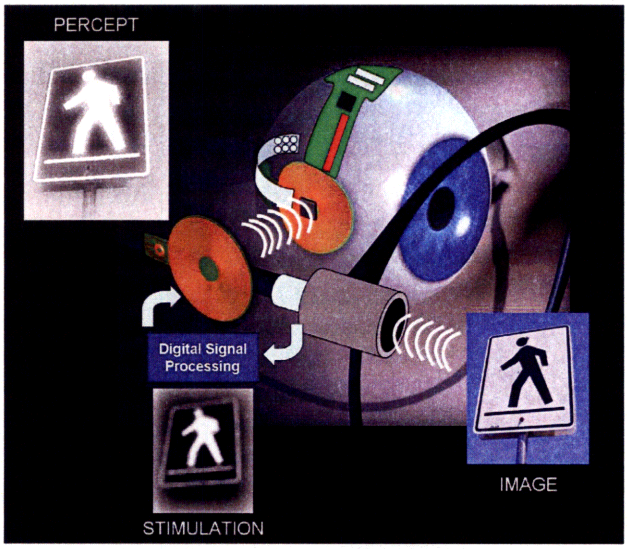

Figure 1-1: A retinal implant. A visual scene is captured by a camera and subsequently analyzed in order to be converted into an appropriate pattern of electrical stimulation. Electrical current passing from individual electrodes (implanted within the retina) stimulate cells in the appropriate areas of the retina corresponding to the features in the visual scene. (Image and text with permission from `http://www.bostonretinalimplant.org`.)

For this thesis, I have developed a statistical inference algorithm for visual decoding and examined how well it captures retinal physiology and visual scene parameters. The outline of this thesis is as follows: Chapter two provides a background on retinal physiology and the neural coding problem. Chapter three presents the concept of receptive field models and point processes, combining the two into a statistical inference framework. Chapter four applies this framework to the problems of estimating neural model parameters and visual stimulus parameters using artificial movies characterized by a small set of global motion parameters. Chapter five goes a step further by focusing on stationary natural stimuli, looking at image identification rather than reconstruction. Chapter six reviews the conclusions of this thesis and provides some direction for future research in this field. Experimental procedures and a more detailed discussion of neural firing rates are presented in the appendices.

Chapter 2

Background

Before addressing the problem of visual neural coding, a good place to start is with a basic understanding of retinal structure and function. This chapter provides a broad overview of the retina and examines existing methods for analyzing neural behavior, with much of the material in the first section adapted from chapter 3 of Hubel[12].

2.1 The Retina

The retina is the part of the inner eye that receives light and converts it to neural signals. In the popular metaphor of treating the eye as a camera, the retina can be thought of as the film. For our purposes, we can think of the retina as consisting of three main neuronal layers, as arranged in Figure 2-1. At the back are the photoreceptors (rods and cones), which convert light into electrical signals. In the middle layer are a collection of bipolar cells, horizontal cells, and amacrine cells. Bipolar cells link receptors directly to ganglion cells, horizontal cells link receptors and bipolar cells, and amacrine cells link bipolar cells and ganglion cells. Finally, at the front of the retina are the ganglion cells, whose axons collect in a bundle to form the optic nerve.

As each retina contains about 125 million rods and cones but only 1 million ganglion cells, it must perform some encoding to preserve visual information along the pathway. Information from the photoreceptors can reach the ganglion cells through

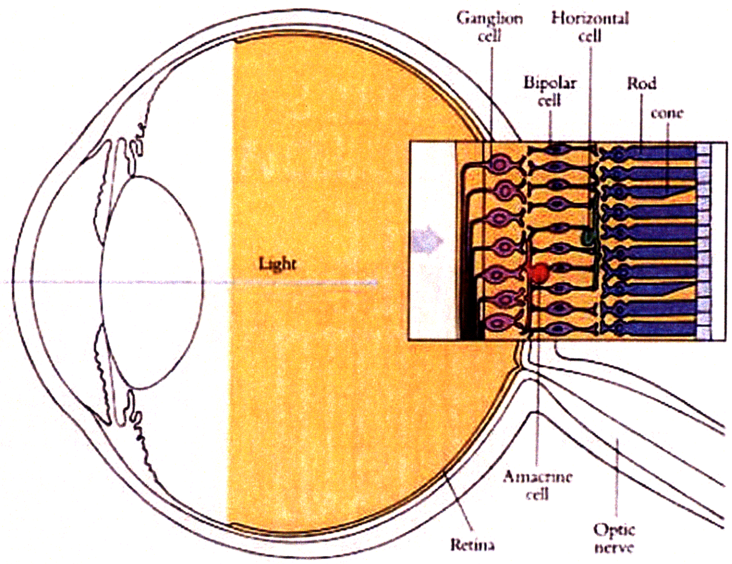

Figure 2-1: The structure of the eye. The enlarged retina at right depicts the retinal layers and cell types. (Image from [12].)

two pathways: directly from receptors to bipolar cells to ganglion cells, and indirectly in which horizontal cells and amacrine cells also modify the signals. The direct path is highly compact, with a few receptors feeding into a bipolar cell, and a few bipolar cells feeding into a ganglion cell. By comparison, the horizontal cells and amacrine cells make wide lateral connections so that receptors spanning a wide area may feed into a horizontal or amacrine cell.

We call the region of space in which the presence of a stimulus alters the firing rate of a neuron its receptive field (RF), and for ganglion cells, this RF roughly corresponds to the area occupied by the receptors feeding into the cell. Experiments have shown that ganglion cell RFs are characterized by center-surround behavior: they consist of a small annular center where light has one effect on cell firing, and a larger annular surround, where light has the opposite effect on cell firing. ON center cells are excited by light in center and inhibited by light in the surround, whereas OFF center cells exhibit the opposite behavior (see Figure 2-2). Rather than tracing the excitatory and inhibitory inputs from the photoreceptors to the ganglion cell, however (a complicated endeavor as the concepts of excitatory and inhibitory sometimes go in conflict with common sense notions of activation and deactivation), we instead focus on the effects of the bipolar, horizontal, and amacrine cells on the ganglion cells.

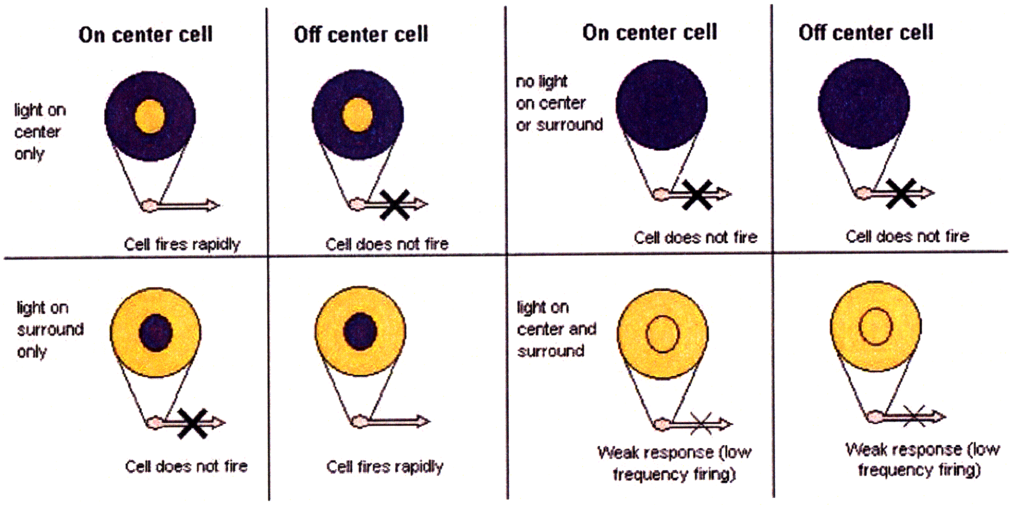

Figure 2-2: Center-surround receptive fields. Response of ON and OFF cells to various light stimuli. (Image from `http://en.wikipedia.org/wiki/Retina`.)

Bipolar cells can be separated into ON center bipolar and OFF center bipolar cells, and the connections of bipolar to ganglion cells are probably all excitatory; thus, ON center bipolars feed ON center ganglion cells and OFF center bipolars feed OFF center ganglion cells. Horizontal cells contribute to the receptive field surrounds of ganglion cells, evidenced by how they affect bipolar cells and by the observation that the area over which receptors feed into a horizontal cell corresponds to the receptive fields of the associated horizontal, bipolar, and ganglion cells. The function of amacrine cells is less understood, and they may or may not take part in the center-surround behavior of ganglion cell receptive fields.

Understanding the functionality and interconnections of retinal cells would be essential if we were attempting to mimic a neural circuit; for neural modeling, such knowledge motivates our choice of cell models and stimuli. For example, one important consequence of center-surround receptive fields is that the eye responds to relative light intensities rather than absolute intensities, so any stimuli that we present should be tuned so that the image contrast can adequately excite ganglion cells. As another example, links between cells provide clues on the aggregate behavior of cell populations and play a part in cell synchrony.

Before moving on, we make one final observation: our understanding of retinal anatomy far exceeds our understanding of retinal function, which far exceeds our understanding of retinal coding. That is, we do not even know what features of a cell response encode useful stimulus information. Arguments have been made for the total number of spikes, the interspike intervals (ISIs), and the absolute spike timings, among others. Similarly, we have limited knowledge on what features of the stimulus, e.g. luminance, contrast, orientation, spatial frequency, are most salient to the eye.

2.2 Survey of Current Literature

Studies on understanding neural behavior have generally approached the problem through a two-stage process of encoding and decoding. In the former, a neural model is typically hypothesized, and given a visual stimulus, we determine the response. In the latter, the problem is analyzed from the opposite perspective: given a neural response, we determine the visual stimulus that produced it. Furthermore, metrics exist to assess the quality of the resulting response or reconstructed stimulus. For example, spike train metrics have been proposed to quantify the similarity and dissimilarity of spike trains[29], and we can use metrics from the signal processing community to quantify reconstruction accuracy.

At the core of both approaches is the development of a proper neural model and assessing its validity. Such investigations have attracted the interest of many cognitive scientists and neurophysiologists, with studies focusing on developing models grounded in physiology and supported by experimental observations. For the visual system, Rodieck[23] proposed a simple spatiotemporal model for ganglion cell neural firing, with later studies looking at cells along the entire visual pathway and suggesting more complicated receptive field shapes[24, 16, 18] and spiking processes[1, 3]. One of the most well-known models is the linear-nonlinear-Poisson (LNP) model[20, 4], characterized by a linear filter followed by a point linearity followed by Poisson spike generation.

In the field of visual neural decoding, initial studies focused on characterizing the response of single neurons, with more recent studies investigating ensemble responses. Warland et al.[30] used optimized (minimum mean-square error) linear filters to decode spike trains from a population of RGCs stimulated with full-field broadband flicker and found that most of the information present in the stimulus can be extracted through linear operations on the responses. Stanley et al.[25] used a similar approach from responses in the lateral geniculate nucleus to reconstruct natural scenes. Frechette et al.[7] estimated from ganglion cells the speed of moving bars in a known direction by finding the peak response of a collection of cross-correlation based detectors. Jazayeri and Movshon[13] developed a likelihood-based approach by computing a weighted sum of neural tuning curves. Guillory et al.[9] determined the typical firing rate profile of ganglion cells to full-field stimuli of different colors using Peri-Stimulus Histograms (PSTHs) smoothed with a Gaussian kernel, then found the likelihood that a response was evoked by a particular stimulus. Pillow et al.[21] integrated receptive field models with likelihood estimation techniques to predict and decode neural responses from retinal ganglion cells.

This thesis attempts to bridge some of these different methods and restricts itself strictly to the problem of retinal coding. However, we do not claim that the brain uses similar algorithms to those proposed.

Chapter 3

Model-Based Statistical Inference

In this thesis, we focus on two particular instances of the neural decoding problems: using the responses from a collection of ON and OFF RGCs simultaneously recorded from a MEA, estimate the speed and direction of a moving edge of light (see Figure 3-1), or identify a natural scene from a set of possible candidates. We develop a neural model based on physiological characteristics, with model parameters fitted by minimizing a cost function over a training set of spike responses. The stimulus is then estimated using these parameters and a distinct test set of responses, with the accuracy of our algorithm determined by looking at the errors in speed and angle estimation or the errors in scene identification.

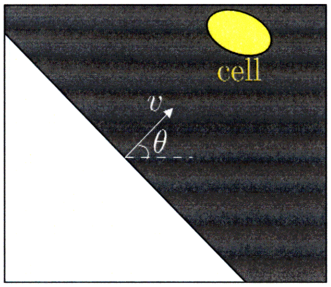

Figure 3-1: Moving edge stimulus. For an ON stimulus, a bright bar of constant intensity moves at a constant speed v and in a constant direction θ. An OFF stimulus is identical except the bright and dark pixels are reversed.

3.1 Firing Rate

We can view a neural response to a visual stimulus through its raster and peri-stimulus time histogram (PSTH) plots. (Sample raster and PSTH plots to moving edges are presented in Figure 4-4.) The raster displays the spike times of a single neuron for repeated trials of the same stimulus and can tell us whether the stimulus resulted in a consistent firing pattern. The PSTH averages the response across these repeated trials to show how the firing rate varies over time and is one way of showing the prototypical response of the cell to the stimulus. Features in the PSTH, such as the maximum firing rate and the time at which it is achieved, or the rate of change in firing rate, or the cross-correlation between the PSTH of two cells, can be used to determine stimulus features.

One way of estimating the visual stimulus when presented with a spike response is to measure how similar the response is to the expected response. Given the spike times $\{t_k\}$ of a cell presented with an unknown visual stimulus, we are thus interested in two components:

1. $\lambda(t)$: the expected time-varying firing rate of the cell when presented with some known stimulus

2. $C(\lambda(t), \{t_k\})$: a cost function measuring the dissimilarity between the expected firing rate and the actual spike times

The estimated visual stimulus is the one that, out of all possible visual stimuli, minimizes $C(\lambda(t), \{t_k\})$.

3.2 Neural Model

We first discuss the expected firing rate $\lambda(t)$. Note that a direct way of computing this function would be to generate a smoothed PSTH for each visual stimulus, but this is impossible to generalize to the case of an arbitrary stimulus, as we would be required to map the firing rate for the infinite range of possible visual stimuli. This direct approach works best as an identification tool, in which we are choosing among a discrete set of known stimuli, but in our definition of the decoding problem, we are interested in a general method that is applicable to a continuous range of stimuli. As such, we instead utilize a model-based approach in which the model is capable of translating any visual stimulus to a firing rate.

A popular model for visual neural firing is the Linear-Nonlinear Poisson (LNP) model depicted in Figure 3-2. A simple interpretation of this model[20] is that the stimulus is linearly filtered by the neuron's spatiotemporal receptive field to produce an intracellular voltage (also known as a generator signal). The voltage is converted via a point nonlinearity to an instantaneous spike rate, and this rate yields a set of spikes via an inhomogeneous Poisson process. We will leave discussion of the Poisson process to the next section when we formulate a cost function; for now, we focus on the linear filter and point nonlinearity. Note that the LNP model is particularly suited to our problem as an inherent assumption is that the neural firing rate captures the essential stimulus information.

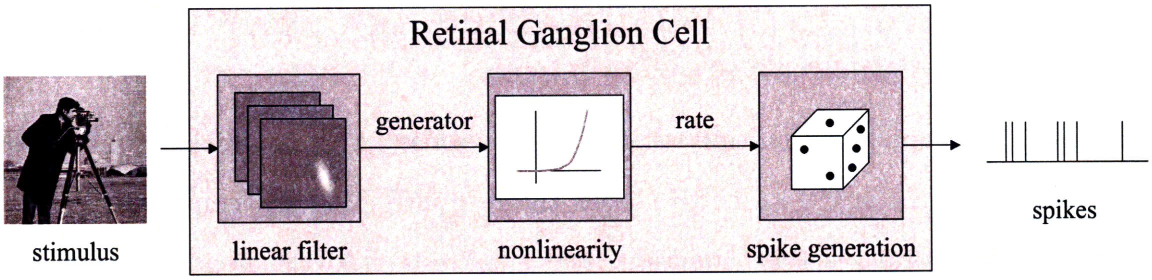

Figure 3-2: Linear-nonlinear Poisson model. The LNP model depicts how a visual stimulus is transformed by the retina into a spike response. The model consists of a linear filter followed by a point nonlinearity followed by Poisson spike generation.

We view the process of generating a rate from a stimulus in terms of the cell's spatiotemporal response. That is, under the assumptions that the cell follows superposition (a cell presented with two images simultaneously produces a response equal to the sum of the responses when presented with the images separately), and that a time shift in the input produces an equal time shift in the output, the cell can be treated as a classical linear, time-invariant system completely characterized by its spatiotemporal impulse response. This is followed by a static nonlinearity to account for any deviations from a strictly linear model, for example due to current thresholds or response saturation. Mathematically, the time-varying firing rate of a cell is given by

$$\lambda(t) = n\left(\int_{-\infty}^{+\infty}\int_{-\infty}^{+\infty} k(x,y,t) * s(x,y,t)dxdy\right), \tag{3.1}$$

where the inner term specifies the generator signal, $k(x,y,t)$ specifies the linear component of the spatiotemporal impulse response (i.e. the spatiotemporal receptive field), $s(x,y,t)$ specifies the stimulus, $n(\cdot)$ specifies an arbitrary point nonlinearity, and the convolution operator $*$ acts in the time dimension only. We make the further simplification that the impulse response is separable such that $k(x,y,t) = f(x,y)\cdot h(t)$, where $f(x,y)$ is the spatial sensitivity function and captures the receptive field shape and size and $h(t)$ is the temporal sensitivity function and captures the impulse response in the center of the receptive field. Though in general, $k(x,y,t)$ is not spatiotemporally separable[22], such an assumption is widely used and greatly reduces model complexity. Thus,

$$\lambda(t) = n\left(\int_{-\infty}^{+\infty}\int_{-\infty}^{+\infty} f(x,y)s(x,y,t)dxdy * h(t)\right). \tag{3.2}$$

The idea is simple: we multiply the spatial response function by the stimulus and integrate over the entire space to obtain the response intensity, convolve the result with the temporal response function, and substitute into the nonlinearity. (Note that some treatments set the firing rate as the convolution of the stimulus with the spatiotemporal receptive field, where convolution acts on both space and time. This is equivalent to our derivation if we flip the spatial response across both axes.)

Numerous methods exist for estimating f, h, and g: using spots of light of increasing diameter, drifting bars or gratings, or white noise analysis[22, 4]. The last of these is also known as the reverse correlation or spike-triggered average (STA) approach, and we choose this method as it is highly robust and easily scalable to multi-cell recordings. In this approach, the retina is presented with spatiotemporal white noise, and the average stimulus preceding a spike is determined. This mean effective stimulus or spike-triggered average directly maps to $f(x,y)$ and $h(t)$. The nonlinearity $n(\cdot)$ can then be estimated by binning the response and plotting the average spike count in bins with nearly equal linear response components.

To complete the model, we choose simple functions for the spatial and temporal sensitivity functions and for the nonlinearity. These functions have a basis in experimental observations and by choosing simple functions, we make the model analytically tractable while reducing the risk of overfitting. We arbitrarily scale the spatial and temporal sensitivity functions, with the relative amplitude of the firing rate across cells being determined mainly through the nonlinearity.

We take a moment to note that this model is highly generalizable and allows us to capture RGC behavior for an arbitrary stimulus. In practice, however, we desire closed form solutions to the convolution formula, which is only possible with intelligent choices for the various cell model components and for simple classes of stimuli such as the moving edge stimulus.

3.2.1 Spatial Sensitivity

The standard spatial structure for ON and OFF RGC receptive fields consists of a difference of Gaussians (with common mean and variance shape but differing amplitudes and variance sizes) to capture their center-excited, surround-inhibited behavior[23, 5]:

$$f(\underline{x}) = \pm \left(\frac{1}{2\pi|\mathbf{\Sigma}|^{1/2}} \, e^{-\frac{1}{2}\left(\underline{x}-\underline{\mu}\right)^T \mathbf{\Sigma}^{-1}\left(\underline{x}-\underline{\mu}\right)} - k\frac{1}{2\pi|r\mathbf{\Sigma}r|^{1/2}} \, e^{-\frac{1}{2}\left(\underline{x}-\underline{\mu}\right)^T (r\mathbf{\Sigma}r)^{-1}\left(\underline{x}-\underline{\mu}\right)} \right), \tag{3.3}$$

where $\underline{x} = [x, y]^T$ specifies the spatial location, $\underline{\mu} = [\mu_x, \mu_y]^T$ specifies the center of the RF, $\mathbf{\Sigma} = \begin{bmatrix} \sigma_x^2 & \sigma_{xy} \\ \sigma_{xy} & \sigma_y^2 \end{bmatrix}$ specifies the covariance (size, shape, and rotation) of the RF center, k specifies the relative strength of the surround, r specifies the relative size of the surround, and the initial sign denotes the cell polarity (+ for ON cells and − for OFF cells). To enforce center-surround behavior, we must have $0 < k < 1$ and $r > 1$. We note that a Gaussian is particularly useful from a mathematical standpoint as it is invariant to rotation. That is, we can transform $f(x, y)$ above to a new coordinate system $\underline{\xi} = [\xi, \eta]^T$, and the resulting spatial function $f'(\xi, \eta)$ is still a Gaussian, though with a translated and rotated mean and covariance matrix. This trick will be useful in determining the rate function for a moving edge in Section 3.2.5.

3.2.2 Temporal Sensitivity

The temporal sensitivity function captures the memory of a neuron and accounts for any delays in the response or for a response that decays with time. If we view the time-reversed impulse response, we can also treat the temporal response as capturing how strongly the stimulus at a previous time affects the present firing rate. Its structure depends on the RGC type (e.g. brisk, transient, delayed, etc), but experiments have shown that many cells display a biphasic profile that can be described as a difference of two temporal low-pass filters[5, 6]:

$$h(t) = \left[p_1(t/\tau_1)e^{-t/\tau_1} - p_2(t/\tau_2)e^{-t/\tau_2}\right] u(t), \tag{3.4}$$

where t specifies the time after the present, p_1 and p_2 specify the relative strengths of the positive and negative lobes, τ_1 and τ_2 specify the decay rate of the two lobes, and $u(t)$ represents the unit step and ensures that the temporal response is causal. In this formulation, all parameters should only take on positive values.

While this equation more accurately reflects the shape of temporal sensitivity function, it is not conducive to a closed form convolution formula using a moving edge stimulus. Therefore, for simplicity, we instead use a modified version of the decaying exponential *step response* proposed by Rodieck[23] that captures transient and maintained characteristics with a monophasic profile:

$$h_s(t) = \alpha^2 \left[(t-\tau)e^{-\alpha(t-\tau)} + m\right] u(t-\tau), \tag{3.5}$$

where t and $u(t)$ are as above, τ specifies the minimum delay, α specifies the reciprocal time constant in the transient decay, and m specifies the maintained firing rate. Similar to above, α is strictly positive, and τ and m are nonnegative. Note that the derivative of this step response would display a similar biphasic profile to that of Equation (3.4). Furthermore, we have normalized the transient term of $h_s(t)$ so that we can interpret it as the probability density function of time delay, i.e. the delay in response follows a random distribution according to $h_s(t)$. This comes in part from generalizing $h_s(t)$ from a delayed impulse: rather than all of the response from a stimulus falling at a single time indicated by the impulse, it is spread according to $h_s(t)$. We note that this normalization is introduced solely for interpretative reasons; it is in general unnecessary since similar scalar changes can be introduced in the spatial or nonlinearity functions.

3.2.3 Nonlinearity

The nonlinearity seems to play a less important role than the spatiotemporal receptive field in the case of retinal modeling as an approximately linear relationship exists between luminance and the firing activity of retinal cells. However, experiments have shown the existence of a nonlinearity which can be captured using the lower portion of a sigmoidal function[23, 5]:

$$n(g) = a\,\Phi(bg - c) + d, \tag{3.6}$$

where g specifies the generator signal, $n(g)$ specifies the firing rate, $\Phi(\,\cdot\,)$ specifies the cumulative standard normal distribution function, and a, b, c, d are free parameters that specify the shape of the nonlinearity, where $a, b, c > 0$. To provide some intuition, a specifies the firing rate intensity, b and c account for thresholding or saturation, and d specifies the spontaneous background rate.

As with the temporal sensitivity function, we forego this more accurate definition in favor of a simple *max* function

$$n(x) = \max(ax, d) \tag{3.7}$$

.

3.2.4 Stimulus

We have been somewhat lax in our definition of stimulus thus far, using it to refer to whatever is presented to the retina during some time period. However, a stimulus can be represented by its luminance, contrast, or some other feature. Given our knowledge that the eye is sensitive to relative intensities, it seems reasonable to let $s(x, y, t)$ capture contrast (as measured in deviations from a mean intensity).

For the case of a binary stimulus, the parametrization of the visual stimulus is somewhat trivial. Following the plus-minus convention set out in the spatial sensitivity function, we let $+1$ and -1 represent the brightest and darkest intensities, respectively. For a moving edge, we say that an ON stimulus represents a dark-to-bright transition, and an OFF stimulus represents a bright-to-dark transition. For natural scenes, we use the 8-bit pixel values $[0, 255]$ rescaled to lie within $[-1, +1]$.

3.2.5 Rate Function

Due to the center-surround characteristics of the RF, a moving edge stimulus elicits a strong response in cells of matching contrast polarity (ON/OFF) and a weak response in cells of opposite polarity. However, we look only at the case where the stimulus and RGC have matching polarity. Then, putting the above components together, the generator signal in response to a moving edge of (v, θ) is given by

$$g(t) = (g_{ex,m} + g_{ex,t}) - (g_{in,m} + g_{in,t}), \tag{3.8}$$

where

$$\begin{aligned}
g_{ex,m} &= m\,\alpha^2\,\Phi(x'; \mu, \sigma) \\
g_{ex,t} &= \tfrac{\alpha^2}{v}\, e^{-\frac{\alpha}{v}[x' - (\mu + z_{ex}/2)]}\, \left[(x' - (\mu + z_{ex}))\,\Phi(x'; \mu + z_{ex}, \sigma) + \sigma^2\, N(x'; \mu + z_{ex}, \sigma)\right] \\
g_{in,m} &= k\, m\,\alpha^2\,\Phi(x'; \mu, r\sigma) \\
g_{in,t} &= k\,\tfrac{\alpha^2}{v}\, e^{-\frac{\alpha}{v}[x' - (\mu + z_{in}/2)]}\, \left[(x' - (\mu + z_{in}))\,\Phi(x'; \mu + z_{in}, r\sigma) + \sigma^2\, N(x'; \mu + z_{in}, r\sigma)\right]
\end{aligned}$$

and the subscripts *ex*, *in*, *m*, and *t* represent excitatory, inhibitory, maintained, and transient components, respectively, and we have used the definitions

$$\begin{aligned}
\mu &= \mu_x \cos(\theta) + \mu_y \sin(\theta) \\
\sigma^2 &= \sigma_x^2 \cos^2(\theta) + 2\sigma_{xy} \cos(\theta)\sin(\theta) + \sigma_y^2 \sin^2(\theta) \\
x' &= v(t-\tau) \\
z_{ex} &= \sigma^2 \alpha / v \\
z_{in} &= (r\sigma)^2 \alpha / v,
\end{aligned}$$

where $N(x;\mu,\sigma)$ specifies the probability density function of a normally distributed random variable with mean μ and standard deviation σ, and $\Phi(x;\mu,\sigma)$ specifies its associated cumulative distribution function. As expected, for small v, the response approaches the difference of Gaussians indicated by $f(x,y)$, and for large v, the response approaches the decaying exponential indicated by $h(t)$. See Appendix B for the derivation and further discussion of this generator signal.

For an arbitrary stimuli such as a natural scene, no closed form solution exists for the firing rate function. Instead, we must perform convolution (or convert to the Fourier domain and use the FFT) to determine the generator signal. For analytic tractability, we restrict ourselves to stationary images so that only a single spatial summation is required per image.

3.3 Cost Functions

There exist numerous metrics for comparing spike trains[29], many of which trade-off different assumptions. We limit our attention to some of the most popular choices.

3.3.1 Maximum Likelihood

In the maximum likelihood (ML) approach, we choose a cost function that incorporates the above receptive field model into a statistical inference framework. That is, given the firing rate of a cell as determined by its model parameters, we can treat the spike times as a point process and determine the probability of a sequence of spike events. Finding the stimulus parameters then becomes a problem of maximizing the likelihood of the spike response. To derive the likelihood formula, we make use of the Poisson assumption in the LNP model. For a cell that fires according to an inhomogeneous Poisson process with rate $\lambda(t)$, we can easily compute the likelihood of K spikes occurring at $X = \{t_k\}_{k=0}^{K-1}$ in the stimulus interval $(0, T]$ since a Poisson process is memoryless with independent intervals between spike events:

$$\begin{aligned} f(X) &= f(t_0, t_1, \ldots, t_{K-1}, t_K > T) \\ &= f(x_0 = t_0, x_1 = t_1 - t_0, \ldots, x_{K-1} = t_{K-1} - t_{K-2}, x_K > T - t_{K-1}) \\ &= \lambda(t_0) \exp\left(-\int_0^{t_0} \lambda(\tau) d\tau\right) \cdot \lambda(t_1) \exp\left(-\int_{t_0}^{t_1} \lambda(\tau) d\tau\right) \ldots \\ &\qquad \cdot \lambda(t_{K-1}) \exp\left(-\int_{t_{K-2}}^{t_{K-1}} \lambda(\tau) d\tau\right) \cdot \exp\left(-\int_{t_{K-1}}^{T} \lambda(\tau) d\tau\right) \\ &= \left(\prod_{k=0}^{K-1} \lambda(t_k)\right) \left[\exp\left(-\int_0^T \lambda(\tau) d\tau\right)\right] \end{aligned} \tag{3.9}$$

or taking the log likelihood,

$$L(X) = \ln(f(X)) = \sum_{k=0}^{K-1} \ln \lambda(t_k) - \int_0^T \lambda(\tau) d\tau. \tag{3.10}$$

Note that Equation (3.9) is exactly the probability density for the set of spikes X.

In analyzing Equation (3.10), we see that the two constituent terms act in opposing manners. That is, the summation term acts on the actual spike times: spikes occurring at probable times increase the log likelihood and spikes occurring at improbable times decrease the log likelihood. At the same time, the integral term limits the overall firing rate. Therefore, a consistently high rate is favorable to the first term at the cost of the second term; similarly, a consistently low rate is favorable to the second term at the cost of the first term. This leads to the rather intuitive conclusion that given a set of spikes, firing rate that achieves a low-cost should only attain non-zero values when spikes are present (due to the first term) while being near-zero elsewhere (due to the second term).

We should also remember that any neural response will typically include some spontaneous spikes that contribute no useful information about the projected stimulus. However, the summation term treats stimuli-triggered spikes and spontaneous spikes equally. We postulate that the brain can determine whether a spike contains useful information or whether it can be treated as a statistical outlier. With this hypothesis, we arrive at the modified likelihood formulation

$$f_w(X) = \left(\prod_{k=0}^{K-1} w(t_k)\lambda(t_k)\right)\left[\exp\left(-\int_0^T \lambda(\tau)d\tau\right)\right] \tag{3.11}$$

$$L_w(X) = \sum_{k=0}^{K-1} w(t_k)\ln\lambda(t_k) - \int_0^T \lambda(\tau)d\tau, \tag{3.12}$$

where $w(t_k)$ specifies a weighting function. For example, to simply ignore spikes with low likelihoods, we could use

$$w(t_k) = \begin{cases} 0, & \lambda(t_k) < P \\ 1, & \lambda(t_k) \geq P, \end{cases}$$

where P represents a percentile, e.g. $P = 0.05$ indicates we would ignore spikes with likelihoods in the lowest 5%. As an alternative, we could consider the function $w(t_k) = w$, where $w > 1$, so that having a firing rate that reflects the spike times is more important than maintaining a minimal rate.

When using more than one cell response, we make a further simplification that the cells act independently so that we can take the joint likelihood of a spike response as simply the product of the individual likelihoods. While such an assumption is not true in general (for a simple counterexample, consider two ganglion cells that share inputs from common amacrine and horizontal cells), allowing for dependent spike generation requires more complicated analysis in order to capture the network dependencies or requires large amounts of training data to factor the dependencies into the individual cell models. With the assumption of independence then, for N cells with rates $\lambda_n(t)$ and spikes X_n,

$$f(\{X_n\}_{n=0}^{N-1}) = \prod_{n=0}^{N-1} f_w(X_n) \tag{3.13}$$

$$L(\{X_n\}_{n=0}^{N-1}) = \sum_{n=0}^{N-1} L_w(X_n). \tag{3.14}$$

3.3.2 Metrics Using Temporal Smoothing

If we wish to compare rate functions directly, we can use other measures such as correlation or mean integrated square error (MISE $= \frac{1}{T}\int_T (\lambda(t) - \tilde{\lambda}(t))^2 dt$), both of which remove the Poisson assumption on the firing rate process, with the tradeoff that they require statistical smoothing of the spike response. In this approach, the experimental spike train is converted to a firing rate $\tilde{\lambda}(t)$ by applying a smoothing kernel, and the optimal model or stimuli parameters are chosen to maximize the correlation or minimize the MISE between the expected and experimental firing rates. One can think of applying the smoothing kernel as a method converting the PSTH into a continuous time function, where in this case, the kernel changes each spike instance, as opposed to each spike bin, into a distribution. This temporal filtering therefore allows for the stochastic variability inherent in spike activity. Of course, requiring a filter means that the kernel characteristics become another parameter in our model, but most research has found that a Gaussian filter with a fixed width of 10 ms is optimal[7],[9].

3.4 Summary and Contributions

In this chapter, we have developed a method for comparing observed spike responses to expected firing rates. In doing so, we presented the linear-nonlinear Poisson model for characterizing RGC neural behavior, and we gave examples of spatial, temporal, and nonlinearity response functions that are both analytically tractable and supported by physiological observations. We also introduced various cost criteria for comparing observed and expected responses. We will combine the LNP model with the ML cost function in the next two chapters and show how they can be applied to estimating neural model and visual stimulus parameters.

Before concluding this chapter, we note that this approach combining neural models with a likelihood cost function is also adopted by Pillow et al.[21], though they restricted their stimuli to full-field pulses and were therefore not able to estimate any spatial characteristics for the cells. Furthermore, rather than perform full reconstruction, they tested the model's decoding capabilities for the simple case of inferring which one of two possible stimuli was most likely given a set of observed neural responses.

This work is also a generalization of the method presented previously by the group (see Chapter 7-8 of [28]). [28] noted that the response of a RGC to a moving edge had a sharp peak when the edge passed over the cell, and thus modeled the RGC RF as a 2D Gaussian. We have extended this RF to the more generalized LNP model and, in Appendix B, shown sufficient conditions on the stimulus and on the spatial, temporal, and nonlinearity functions in order to obtain a closed form rate equation. Furthermore, we have introduced a more robust likelihood formulation and presented the idea of maximum likelihood as that of minimizing a spike metric-based cost criterion. In short, we have extended the work of [21] and [28] to a general framework incorporating neural models and minimum cost optimization.

Chapter 4

Decoding Global Motion

In this chapter, we focus our attention on a moving edge stimulus consisting of a bar of constant intensity moving at a constant speed v and in a constant direction θ, where the edge passes over the center of the screen at time $t = 0$. Analysis was separated into a training phase in which we knew the visual stimulus and examined the RGC outputs to determine the neural model parameters, and a testing phase in which we examined RGC outputs elicited by moving edges and estimated the unknown speeds and directions. In the both cases, the unknown parameters (either for the model or for the stimulus) were chosen to minimize the cost of the observed responses. Moving bars of constant speed, direction, and contrast were presented to the retina, where the speed and directions probed were (300 µm/s, 600 µm/s, 900 µm/s, 1200 µm/s) and (0°, ±90°, 180°), and the contrast was ±25%. Experiments were performed by recording in-vitro responses using multi-electrode arrays with rabbit retina. Data analysis is discussed in this chapter; for more details on the experimental procedure, refer to Appendix A.

In our discussion, a trial refers to a single presentation of a moving edge stimulus over a retina, and a condition refers to a single retina tested with a specific speed, direction, and polarity. In total, we ran 4 retina stimulated with moving bars at 2 polarities, moving at 4 speeds in 4 directions, yielding 128 conditions. Since some of the retina responded only to ON or OFF bars, however, the experiments yielded 96 conditions.

4.1 Data Analysis Procedure

4.1.1 Pre-processing

Before either training or testing, we first visually inspected (1) the STA and autocorrelation in response to the m-sequence and (2) the PSTH and raster plots in response to the moving bars. The STA, PSTH, and raster plots were defined in Chapter 3; the autocorrelation of a spike train plots the the mean firing rate of a cell as a function of time after occurrence of a spike. For the PSTH, we used bins with width equal to the reciprocal of the frame rate. Only cells that exhibited typical STAs and autocorrelation functions[6] and the expected firing pattern to moving bars (e.g. maximal response as the bar passes over the RF, consistent responses across trials) were selected for further analysis. Cells that were rejected tended to display spontaneous firing such that their responses arise from internal factors rather than being a result of external stimuli, or were hyperactive. This selection also reduces the chance of using cells that may have been found through spike sorting errors. Typically, a cell either satisfied or did not satisfy both criteria. A few cells responded well to the m-sequence but not to moving edges; the STAs and autocorrelations of such cells matched those of local edge detectors[6], so this behavior is expected. Others responded well to the moving edges but not to the m-sequence; they generally had small RFs that were poorly captured with our m-sequence resolution. Finally, cells that responded only to a subset of speeds or were directionally-selective were also rejected.

4.1.2 Training

In the training phase, model parameters were found independently for each cell. We initialized our estimates by finding the best-fit STA through least-squares minimization. An m-sequence of 32,768 frames was presented to the retina, and the mean stimulus over the 303 ms preceding a spike was found. This generated roughly 20 frames of 16×16 elements. For each frame, the intensities over all elements was found.

Then the frame with maximum absolute sum was fit to the desired spatial sensitivity function, and the intensity over the 3×3 center pixels of this profile were averaged for each frame to obtain a temporal profile. If a step response was required, the temporal profile was integrated before being fit to the temporal sensitivity function. Finally, the generator signal was obtained using the fitted spatial and temporal profiles and compared to the observed firing rate to obtain the nonlinearity. The optimizations for the spatial, temporal, and nonlinearity parameters were performed numerically using the Gauss-Newton method to minimize the mean squared residual error.

One known problem of STA is that if the pixel values are binary, as in the case of m-sequences, the STA provides a biased estimate of the receptive field[4]. For these reasons, and also to take into account possible behavioral changes due to the moving edge stimuli, the estimates were refined through numerical optimization using sequential quadratic programming to minimize the cost of the observed responses to moving edges. We used the responses from a subset of trials, where we chose only stimuli with a matching polarity but we had the same number of trials for all (v, θ) pairs. (E.g. An ON cell is trained only using ON stimuli, but using equal numbers of responses obtained from (ON, 300 µm/s, 0°) stimuli as from (ON, 600 µm/s, 0°) stimuli, and so on.) Note that we could alternatively have taken the minimum cost estimates over the m-sequence, but in practice, this was computationally prohibitive for optimization.

4.1.3 Testing

In the testing phase, the model parameters were kept constant, and visual parameter estimates were made by minimizing the cost of the observed spike responses. Unlike the training phase, we find the parameters that achieve the minimum collective cost across all cells using the responses from a single trial. We could presumably obtain better results using responses averaged from multiple trials with the same stimulus, but as we will see, even a single trial produces accurate estimates. Furthermore, the brain is not privy to multiple responses during decoding. In this phase, the polarity of the moving edge is assumed to be known, and initial speed and direction estimates for a trial were made by finding the optimal (v, θ) pair using a three-step search method. That is, costs are evaluated at points along a coarse grid, and the $(\hat{v}, \hat{\theta})$ with minimum cost is chosen. This is repeated two more times, each time using a grid of higher resolution centered around the previous $(\hat{v}, \hat{\theta})$. The final estimate was then refined through numerical optimization. To prove the method does not simply learn the training set, the trials used in testing were distinct from those used in training.

4.1.4 Selection of Model and Cost Function

As a reminder, we chose to use a difference of Gaussians spatial profile (7 parameters), a decaying exponential step response temporal profile (3 parameters), and a linear with simple thresholding nonlinearity profile (2 parameters). We focus on the likelihood cost criterion where we ignore spikes with likelihoods in the lowest 5%. While the likelihood formula is not convex so that the optimized estimates may not correspond to the global maximum, the algorithm performed well in practice: model estimates were physiologically meaningful, and stimulus estimates were near the true known values. Also, we note that during training, the second stage of cost minimization using ML is robust to the initial estimates; other ad-hoc methods of estimating the spatial, temporal, and nonlinearity functions, for example, based on measurements of the PSTH response to moving edges, yielded similar optimized parameters, though more iterations were necessary for convergence.

4.2 Results

Data from four rabbit retina were collected, and Tables 4.1 and 4.2 show the number of recorded cells in each retina and their peak firing rates under various stimuli. Note that retinal pieces A and used one setup, and retinas B and D used a second setup. For each retina, the number of cells excludes those that do not display typical STAs or PSTHs from visual inspection. This reduces inflation of cell counts, as visual inspection revealed that some of the units found through spike sorting most likely consisted of spikes leftover from the clustering process. If fewer than five cells of either polarity type were found in a single retina, no analysis was performed for that subset; this accounts for the zero ON cell count for three of the retinas. Anywhere from 36-81% (8/22-22/27) of all recorded cells were used in subsequent analysis. It is possible that the brain would incorporate knowledge from the remaining cells but would preferentially select the chosen cells based on internal metrics of spike behavior. For comparison, prior work with MEAs have typically found around 30-90 usable cells per retina[27, 26]. We postulate that fewer cells were found here due to the broad spatial extent of moving edges. A large bar of high intensity will more often inhibit neural response, as a stimulus with large support that overlaps the center-excited RF will typically also overlap the surround-inhibited RF.

retinal piece	# of cells that responded to m-seq	# of cells that responded to bars	# of cells used in analysis (ON,OFF)
A	25	22	8 (0,8)
B	71	50	33 (14,19)
C	41	49	26 (17,9)
D	31	27	11 (0,11)

Table 4.1: Number of recorded cells responding to various stimuli and selected for data analysis. These numbers exclude those judged from visual inspection to have atypical STAs to m-sequences or atypical PSTHs to moving edges.

retinal piece	background firing rate (Hz)		peak firing rate (Hz)	
	ON	OFF	ON	OFF
A				128.5 ± 35.6
B	2.5 ± 1.4	5.0 ± 4.5	40.1 ± 14.1	47.6 ± 32.8
C			30.3 ± 13.2	54.4 ± 35.3
D		17.7 ± 16.7		45.0 ± 21.3

Table 4.2: Experimentally observed firing rates for cells used in data analysis. For each retina, the background firing rate of the cells during exposure to a spatially uniform dark background, and the peak firing rate elicited by bars moving at 300 μm/s. Peak firing rates were calculated from the PSTH with bin widths equal to the reciprocal of the frame rate. Firing rates are the mean ± SD across the ON or OFF cells used in analysis. Empty cells indicate that either the cells were not subjected to a spatially uniform dark background, or no cells were used in the analysis.

4.2.1 Model Parameters

One method of validating our model is to compare the neural model parameters fit through STA and optimized through ML. This presents the ML algorithm as an alternative to the traditional STA approach for estimating neural response functions. To obtain a single measure of the RF size, we define the RF diameter as twice the geometric mean of the square root of the eigenvalues of Σ, i.e. it is the diameter of the circle with equal area as the 1σ ellipse of the center component of the spatial RF. A direct comparison of the model parameters as in Figure 4-1 may be misleading, however, as there are reasons to expect differing results from the two estimation methods. For example, STA with binary sequences leads to biased estimates of the RF size, and ML using responses to moving edges can confound the spatial and temporal responses. Instead, we note that the distributions of the two measures, after accounting for shifts by subtracting the mean, do not differ significantly (Kolmogorov-Smirnov ON: $k = 0.286$, $n_1 = n_2 = 21$, $p = 0.30$ two-sided, OFF: $k = 0.265$, $n_1 = n_2 = 34$, $p = 0.16$ two-sided). Furthermore, the model parameters fall within accepted ranges[6]. For example, the mean of the ML estimate of the RF diameter across the ON or OFF cells used in analysis for each retina was n/a and 237, 314 and 304, 263 and 268, and n/a and 380 µm.

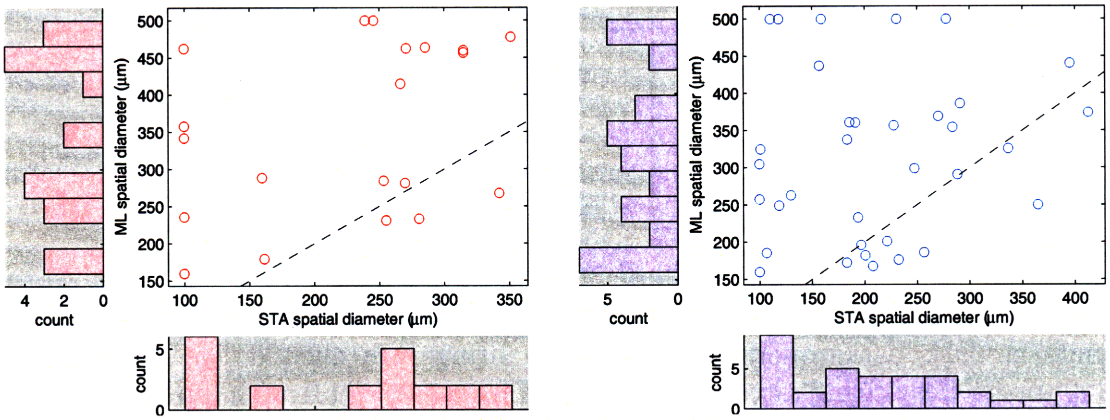

Figure 4-1: Correlation of RF diameters from STA and ML. Data is shown using 21 ON cells (left) and 34 OFF cells (right) from the three retina. The correlation coefficients are 0.43 for ON cells and 0.14 for OFF cells. The dashed line indicates a 1-to-1 fit of the two estimates.

For a single retina, comparison of the RF maps (see Figure 4-2) shows that the optimal RFs found using STA and ML cover roughly the same spatial area. Note again that we do not expect an exact match due to different biases of the two methods. However, we do observe that a cell's RF typically lies in the region surrounding the electrode from which its response was recorded. For a single cell, we can also look at the similarity in the neural response functions, as in Figure 4-3, or the similarity in the experimental and expected firing rates, as in Figure 4-4.

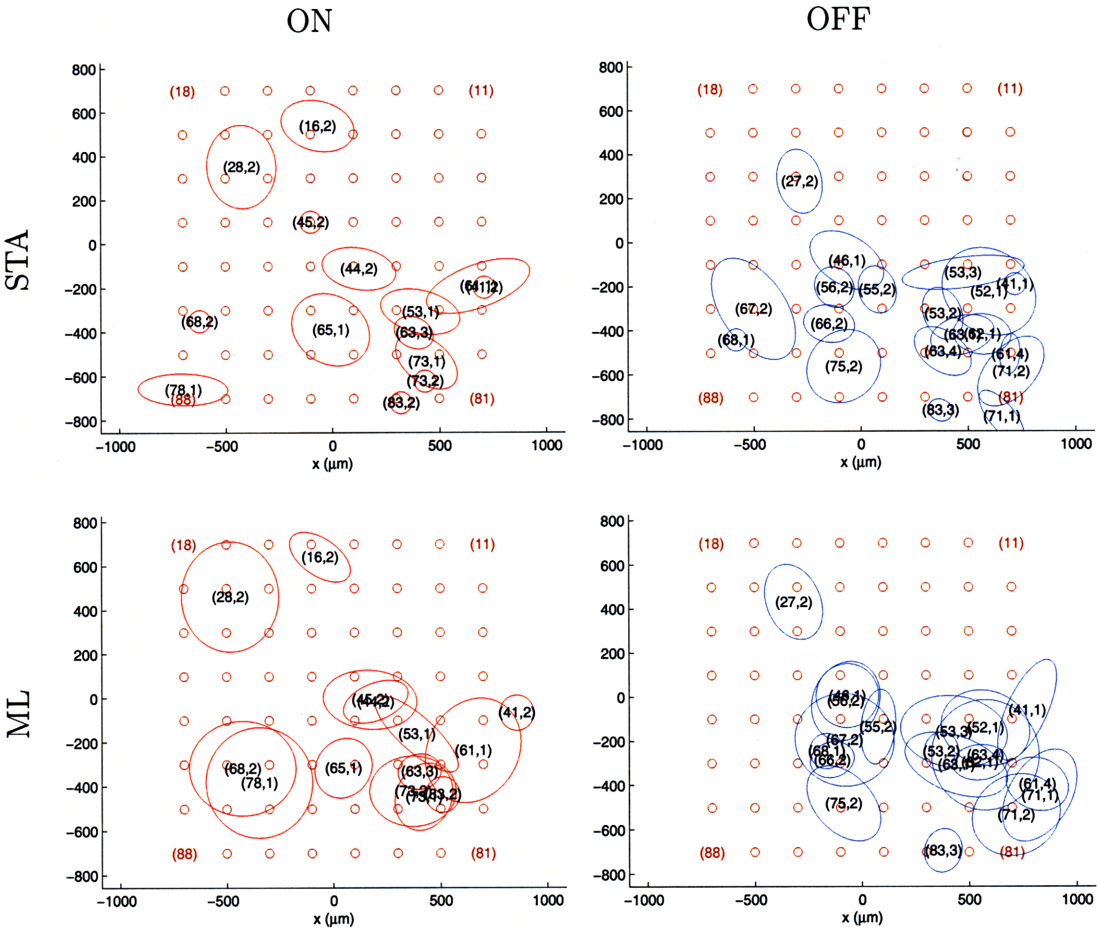

Figure 4-2: RF maps from STA and ML. Spatial RF locations for ON RGCs and OFF RGCs simultaneously recorded in retinal piece B. Outlines represent the 1 SD boundaries of the elliptical Gaussian fits as estimated from STA or ML. Cells are labeled by (electrode,unit), and the brown circles show the location of the electrodes on the MEA for comparison to the cell locations.

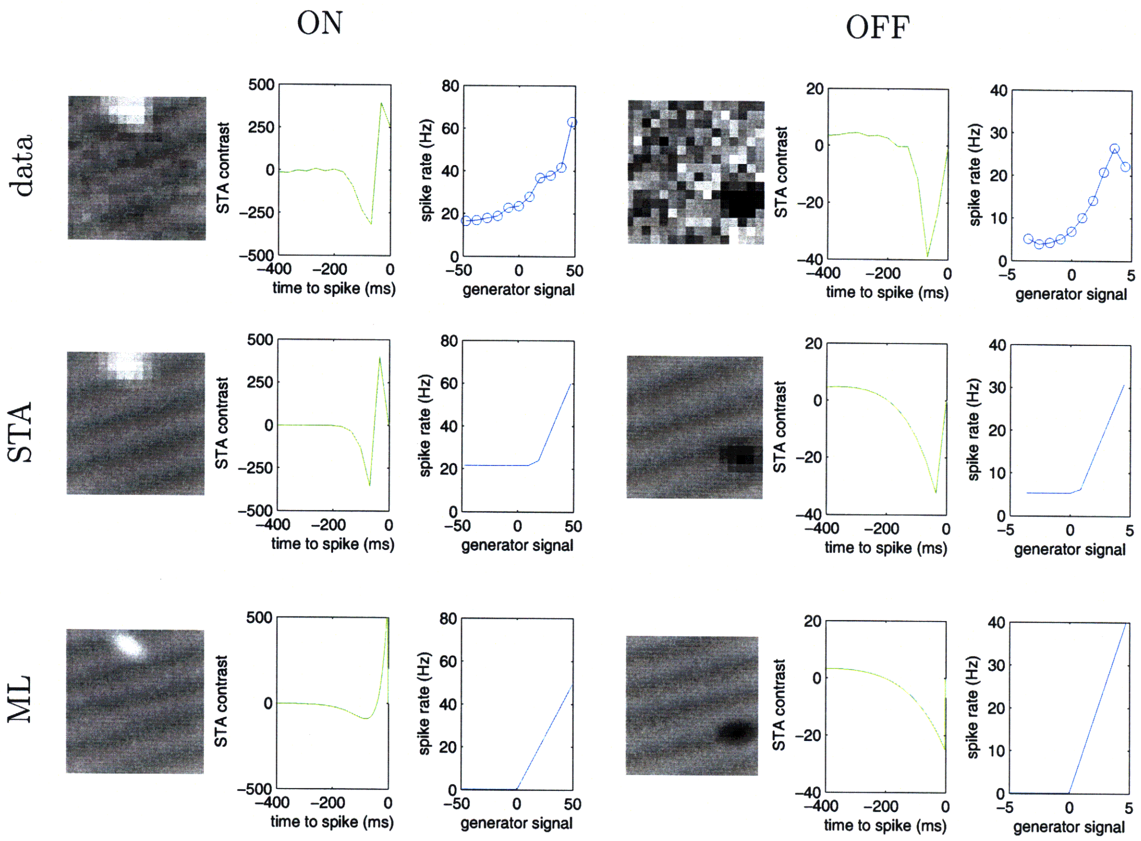

Figure 4-3: Neural response functions and parametric fits. Data is shown for one ON RGC (electrode 16, unit 2 in retinal piece B) and one OFF RGC (electrode 62, unit 1 in retinal piece B). Each set of three panels shows *left*, the average stimulus observed near time-to-peak before a spike (total panel area = $1565\mu\text{m}\times 1565\mu\text{m}$), *middle*, the average time course preceding a spike, summed over the nine center pixels in the center of the RF, and *right*, the average firing rate as a function of the generator signal. These panels map to $f(x, y)$, $h(t)$, and $n(\cdot)$, respectively. From top to bottom, *top*, the panels corresponding to the observed data, *middle*, the parametric fits using STA, and *bottom*, the parametric fits using ML. The difference in the nonlinearity function can be attributed to the cell having different background firing rates during stimulation with a m-sequence and stimulation with moving edges.

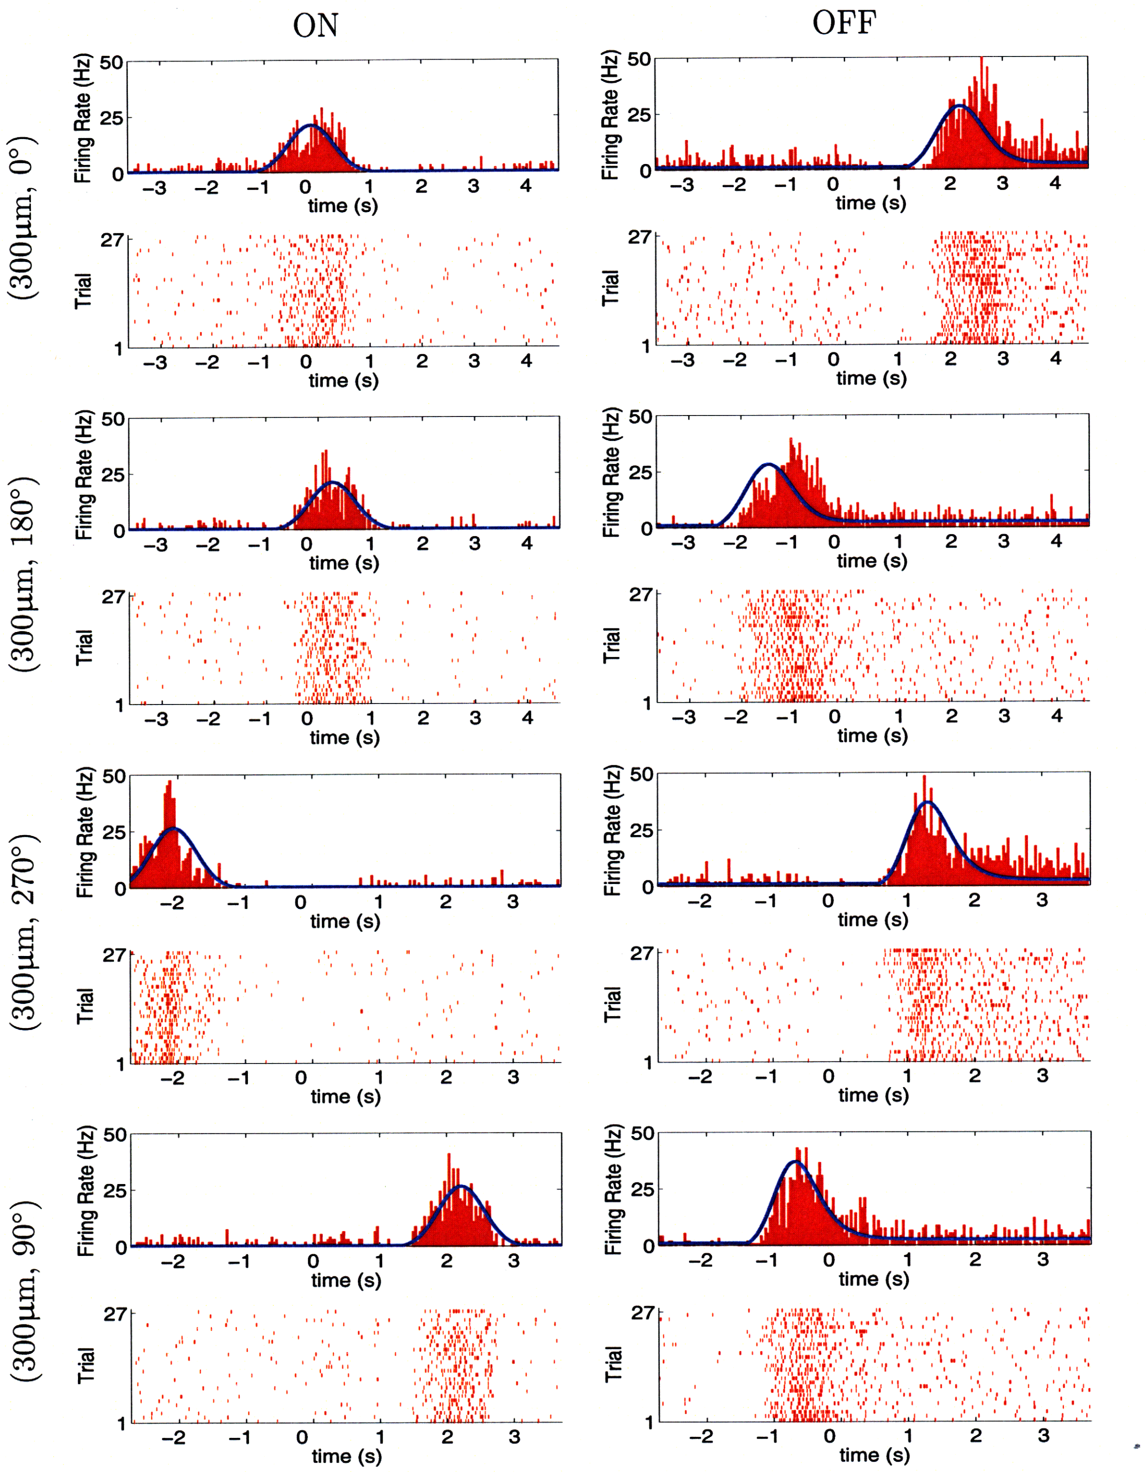

Figure 4-4: Estimated versus expected firing rates. PSTH and raster plots over 27 trials for the ON RGC and OFF RGC of Figure 4-3. From top to bottom, the stimulus was an edge of matching polarity moving at 300 µm/s in the (0°, 180°, 270°, 90°) direction. The model was fit using the data from 48 runs (3 trials, 1 polarity, 4 speeds, 4 directions). The blue line in the PSTH plot indicates the expected firing rate as estimated by the model.

4.2.2 Visual Stimulus Estimates

In this section, we focus on the bias and spread of the visual stimulus estimates as a means of cross-validating our algorithm.

One consequence of the low dimension parameter space for the moving edge is that we can easily visualize the likelihood landscape, as in Figure 4-5. Note that while the landscape may not be convex, there is typically a well-defined global maximum that corresponds to the ML estimate of the visual stimulus parameters. A "ragged" landscape with more local maxima hints at estimates of lower fidelity, as occurred when the responses exhibited larger stochastic variability across trials.

The distribution of speed and direction estimates across all trials for a single condition typically exhibits a Gaussian profile centered near the true stimulus speed and direction (see blue bars of Figure 4-6), so we take the mean and SD as an accurate measure of estimation bias and spread. When looking across various speeds, we use the normalized speed (speed estimate divided by stimulus speed) to reduce the effect of the true speed on our metrics. Figure 4-7 shows histograms of the normalized speed and angle bias and spread for all conditions. The mean normalized speed bias was 0.0017 (SD of 0.0588), indicating a slight tendency to overestimate speed by 0.17%. A direct interpretation of the mean angle bias of -1.04° is not possible since under or overestimation of angle does not have a direct physical meaning when viewed across various true stimulus directions (the most we can say is that a negative bias means the estimate is a clockwise shift of the true direction); instead, we simply note that the angle bias is centered tightly around 0° (SD of 4.58°), indicating little tendency to misestimate the angle. In a similar vein, the distributions of the normalized speed and angle SDs are skewed highly to the right, with peak values of 0.1 and 2.5°, indicating high estimation fidelity.

Simulations

To better understand the limitations of our model, we also conducted simulations of neural responses. To generate a series of spikes from our model, consider that for an

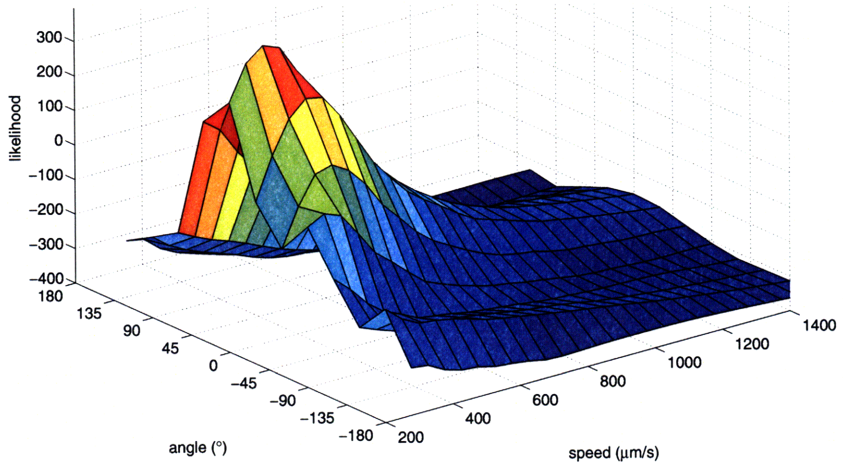

Figure 4-5: Likelihood landscape. The likelihood of (ON, v, θ) stimuli obtained from the responses of the ON cells of retinal piece B in one trial. The ML estimate was found at (305.1 μm/s, 1.7°), compared to the true stimulus of (300 μm/s, 0°).

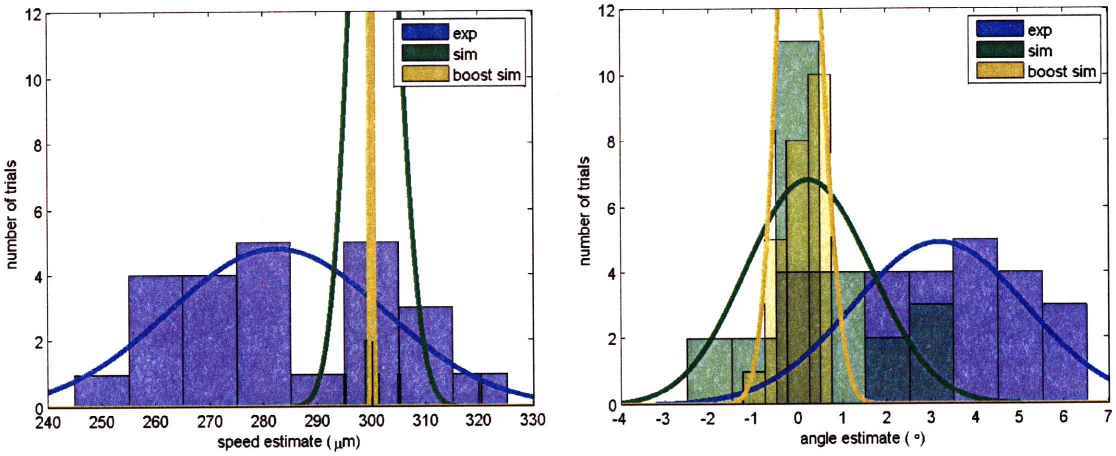

Figure 4-6: Speed and direction estimates. Distribution of speed and direction estimates across 24 repeated trials of (ON, 300 μm/s, 0°) stimuli using the responses of the ON cells of retinal piece B. Blue corresponds to results from experimental responses, green to simulated responses, and yellow to simulated responses with boosted firing rate. The mean ± SD of the speed and direction distributions are 282.5 ± 19.9 μm/s and 3.2 ± 1.9° for experimental responses, 300.7 ± 4.0 μm/s and 0.2 ± 1.4° for simulated responses, and 300.1 ± 0.1 μm/s and 0.06 ± 0.4° for simulated responses with boosted firing rate. Smooth lines are the Gaussian fits with corresponding mean and SD.

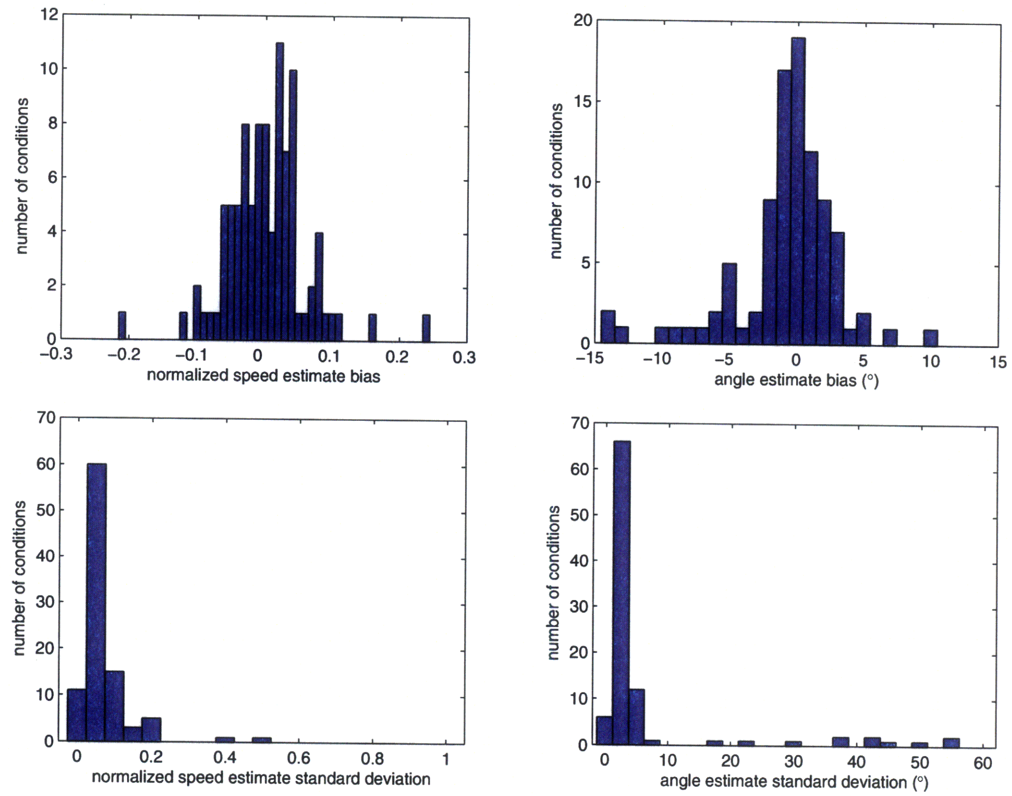

Figure 4-7: Speed and direction estimation bias and variability for experimental responses. Speed estimates are normalized by the true stimulus speed. These values were accumulated across 96 conditions.

inhomogeneous Poisson process with rate $\lambda(t)$, the probability of k events in time bin $[a, b]$ is

$$P(k) = \frac{\exp\left(-\int_a^b \lambda(\tau)d\tau\right)\left(\int_a^b \lambda(\tau)d\tau\right)^k}{k!}. \tag{4.1}$$

In a sufficiently small time bin $[t_0 - \Delta/2, t_0 + \Delta/2]$, the probability of multiple spikes is negligible, so we can approximate a Poisson process as the output of a Bernoulli random variable with $P(1) = \lambda(t_0)\Delta$ and $P(0) = 1 - P(1)$. For our simulations, $\Delta = 0.1$ ms.

In the first study, spike responses were generated using the ML model parameters for the retina. Since the simulated responses follow the proposed model for neural firing, we consider the speed and angle estimates as a baseline against which to compare the estimates obtained from experimental data. However, there is still randomness as the actual spike times are generated through an inhomogeneous Poisson process; thus, we only get a glimpse at the firing rate profile from any given set of spikes. We would expect to obtain a better picture of the firing rate profile if we observe more spike events, so we also artificially boosted the firing rate profile by multiplying the parameters of the nonlinearity by a constant factor, in a sense, increasing the stimuli-triggered firing rate and the background rate of a cell. The green and yellow bars of Figure 4-6 show that distribution of speed and direction estimates using both types of simulated responses (boosting factor = 10), and we see that both bias and spread decrease as we move from experimental to simulated to boosted simulated responses.

One problem with our experimental retina is that we have no control over the cell characteristics, specifically the total number of cells and their spread. To see how these characteristics might affect the fidelity of visual estimates, we constructed "ideal" simulated retina with varying number of cells and varying cellular spread. A simulated retina with parameters (n, d) consisted of n cells hexagonally tiled across a square retinal piece with side length d. All cells within a retina were identical, with model parameters chosen to roughly correspond to the mean or median of those found experimentally. From Figure 4-8, we see that as expected, the accuracy of the speed and direction estimates, as measured by the SD, decreases with the number of cells or as the cells become more clustered.

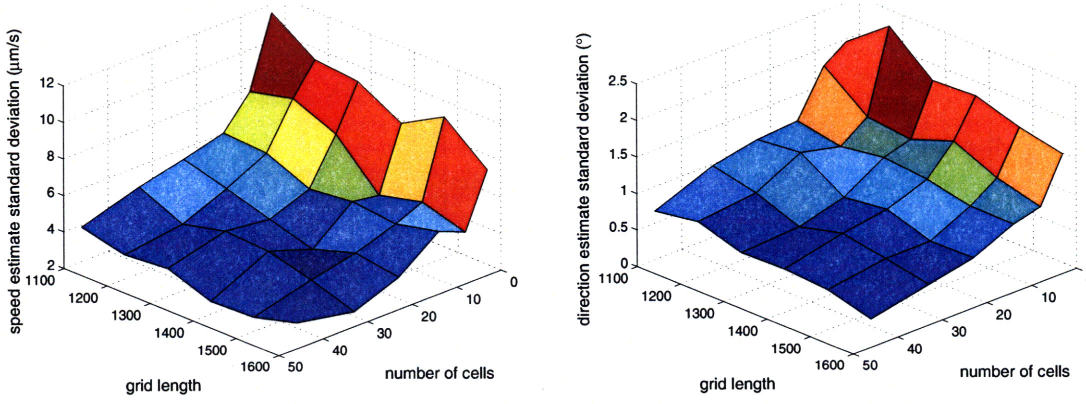

Figure 4-8: Dependence of speed and direction variability on the number n and spread d of cells. Each (n, d) point corresponds to the SD of the estimates using 40 trials of simulated spike responses stimulated with (ON, 300 µm/s, 0°) moving bars.

Comparison to Previous Algorithms

Frechette et al. studied the fidelity of speed estimates across various factors, e.g. neural synchrony and spatial arrangement[7]. We forego a similar in-depth treatment and instead look at how our neural model and ML algorithm compares with his correlation-based approach.

Frechette et al. collected data from 3 macaque monkey retina, with 35-68 ON or OFF parasol cells per retina recorded over a 5° ×10° area, and probed speeds of 7.3°/s, 14.5°/s, 29.0°/s, and 58.1°/s. Compare this to our data from 4 rabbit retina, with 8-17 ON or OFF cells per retina recorded from a 7° ×7° area, and probed speeds of 1.5°/s, 3.0°/s, 4.5°/s, and 6.0°/s. In both animals, we have used the conversion of 200 μm/°. Looking at the speed estimate bias (expressed as a fraction of the SD), Frechette et al. found a mean of -0.3, compared to our mean of 0.0163 (see Figure 4-9a), indicating slight tendencies to underestimate and overestimate speed, respectively. More interestingly, Frechette et al. found that speed variability increases with increasing stimulus speed, but from Figure 4-9b, our ML algorithm shows no such relationship.

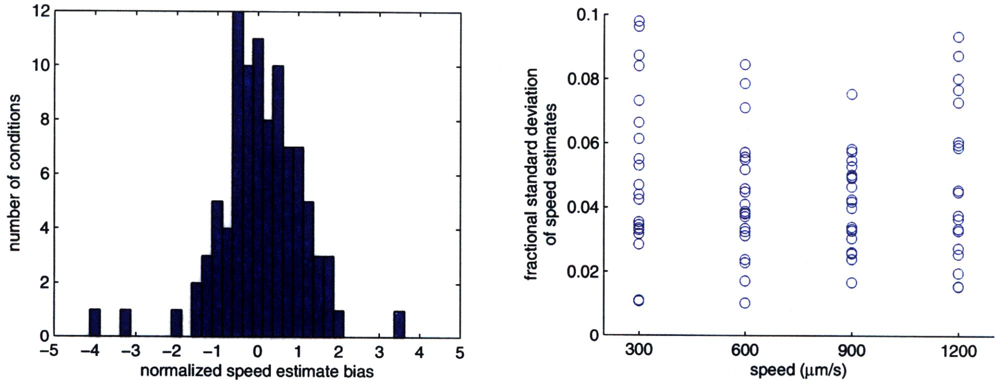

(a) (b)

Figure 4-9: Speed estimate variability and dependence on true speed. These values were accumulated across 96 conditions. *a*, Distribution of speed estimate bias, expressed as a fraction of the speed estimate SD. *b*, Fractional speed estimate variability (SD divided by stimulus speed) as a function of stimulus speed.

4.3 Summary

In this chapter, we have applied the statistical inference algorithm presented in Chapter 3 to the problems of estimating the neural response functions of retinal ganglion cells and estimating the global motion parameters of a moving edge. We have shown that the ML neural parameter estimates are physiologically reasonable and that, when used in the decoding problem, they provide accurate estimates of speed and direction.

Chapter 5

Decoding Natural Scenes

While the model-based statistical inference algorithm is feasible for the simple moving edge of the previous chapter, it becomes intractable when attempting to reconstruct complex scenes. In the fully general discrete case, a movie with $n \times n$ spatial elements and N time frames has $n^2 \cdot N$ parameters over which to optimize. Also, as noted earlier, cells stimulated with arbitrary scenes have no closed form solution for the generator signal of the expected neural response, further making the optimization procedure computationally expensive. Thus, we focus in this chapter on a simplification of the general problem, identification of a stationary natural image. Again, refer to Appendix A for the specifics on the visual protocol.

5.1 Data Analysis Procedure

Due to limited data, cells were chosen based on the pre-processed data of Section 4.1.1, and we used the optimized model parameters over moving edges, as discussed in Section 4.1.2. This introduces many restrictions on the available retinal information. For example, our findings do not make use of local edge detectors (LEDs), which are likely important in natural stimuli encoding[32]. In particular, using only ON and OFF cells means that we focus on the broad spatial features of an image. Furthermore, though we have trained the model parameters using binary stimuli (m-sequences and moving edges), we are attempting to identify novel grayscale images. If our system achieves high identification performance, we therefore have further support that our algorithm generalizes well rather than simply overfitting to the training data.

Similar to before, during testing, the neural model parameters were kept constant, and given these neural parameters and an observed spike response to a test image, the algorithm found the optimal stimulus parameters. However, rather than a continuous optimization problem as for the moving edges of Chapter 4, we now have a discrete identification problem in which the algorithm is told to choose an image out of a set of possible candidate images. That is, for each observed spike response, the likelihood of various candidate images was computed by comparing the experimental response to the expected firing rate evoked by the candidate image. The candidate image whose expected firing rate was most similar to the experimental response, as measured by maximum likelihood, was selected.

5.2 Results

The results in this section come from 10 repeated trials of 10 natural images on retinal piece B. As before, we can compare the experimental and expected firing rates, as in Figure 5-1. Instead of a likelihood landscape over the possible parameters, however, we now obtain a likelihood matrix over the possible image indices, as in Figure 5-2. Interestingly, even though we have ignored LEDs and have trained over binary stimuli, the algorithm is still able to correctly identify grayscale images in many cases. We achieved an identification rate of 50% using the responses from a single trial, and this rate increased to 70% using the collective responses from the 10 repeated trials. Compare this to the 10% expected identification rate for chance performance.

Visual inspection revealed that identification errors tended to occur when the most salient (brightest or darkest pixels) of an image fell within a region outside the 1σ RF of the available cells. In these cases, the image within the region of support of the available cells was unable to produce a sufficiently reliable neural response for comparison to the expected response. Images that were often confused also tended to be visually similar within the spatial RFs of the available cells.

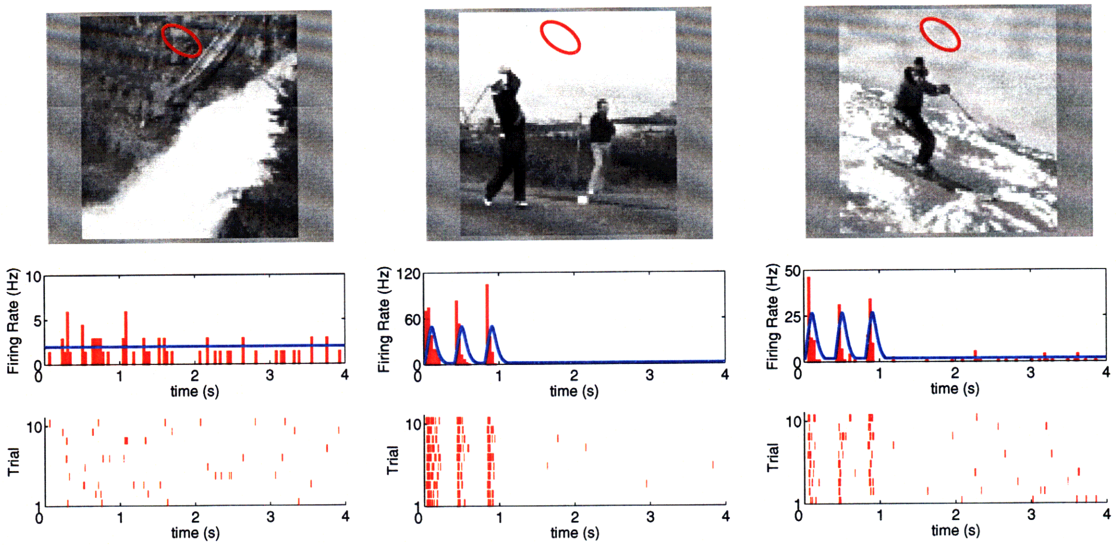

(a) (b) (c)

Figure 5-1: Natural stimuli and estimated versus expected firing rates. Natural images are projected onto the retina in an ON-OFF-ON-OFF-ON pattern, and the PSTH and raster plots are accumulated over 10 trials for the ON RGC whose RF is indicated in red. The blue line in the PSTH plot indicates the expected firing rate as estimated by the model. Note that the cell has a higher firing rate when brighter pixels are projected to its RF, as in *b* and *c*, and that only spontaneous background firing occurs when darker pixels are projected to its RF, as in *a*.

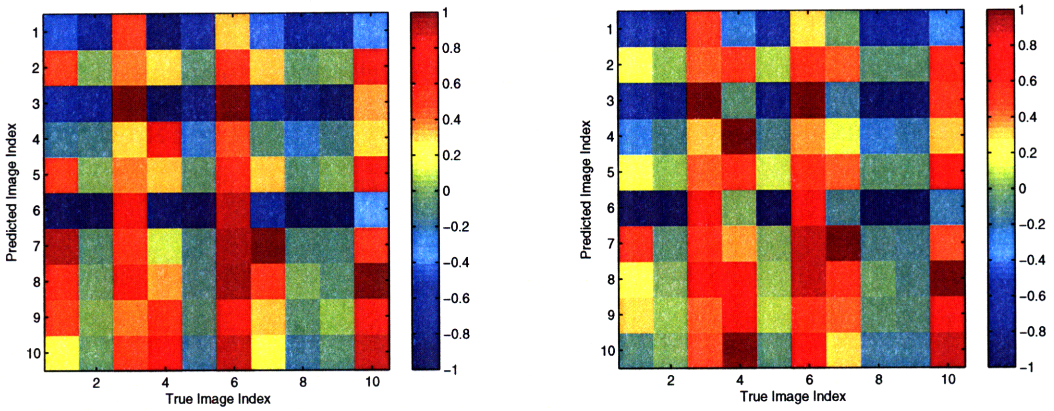

(a) all trials (b) single trial

Figure 5-2: Likelihood matrix for natural images. *a*, The color at the mth column and nth row represents the likelihood of observing the response for the mth image under the prediction that the true image was n. Likelihoods are normalized over all predictions (over each row) and are calculated using all 33 available cells from retinal piece B and all 10 trials. Optimal performance is achieved if the maximum likelihood per column falls along the main (top-left to bottom-right) diagonal. 70% of the images were identified correctly. *b*, Same as *a* using a single trial. 50% of the images were identified correctly.

We also examined how the identification rate changes as we increase the number of candidate images, and as we increased the number of cells (see Figure 5-3). As the set size of candidates images increased to 120 images, accuracy decreased to 20% (repeated trials) and 10% (single trial). Analysis revealed that performance decreases corresponded to adding images with the problems outlined above. In this case, performance no longer changed when increasing the set size past 60 images, suggesting that our model was able to highly accurately predict the responses for a few of the test images, and that these images were sufficiently different from any further introduced candidate images. To explore how identification rate changes with the number of cells, we ranked each cell by its predicative power. Each cell was used in isolation to identify the test images, the likelihoods were normalized to the maximum likelihood the cell achieved over all images, and these normalized likelihoods were summed over all images, with larger total values indicating higher predicative power. Identification performance is not monotonic with the number of cells as this ranking neglects possible correlations across cell predictions; that is, though a cell may perform better than another when viewed in isolation, its contribution to the collective knowledge may be deleterious. However, in general, performance increased with increasing number of cells, with a nearly constant performance achieved after only 12 cells. Analysis revealed that further added cells had RFs that overlapped with previous cells so there was little additional information regarding the input stimulus.

For comparison, Kay et al.[15] used a generative neural network model to predict the fMRI response of the visual cortex. He achieved an identification performance of 92% (12 repeated trials) and 82% (single trial) when identifying a set of 120 images. It is perhaps disingenuous to compare identification performance directly, though, as there are a number of differences in our studies. First, there is a nonlinear relationship between the image luminance and the recorded fMRI activity, and V1 have RFs resembling Gabor wavelets that are also tuned to orientation and spatial frequency. A retinotopy-only model using only the location and size of each voxel's RF performed significantly worse than the Gabor pyramid model, though it still surpassed our achievable rates. Second, their study trained over a set of natural

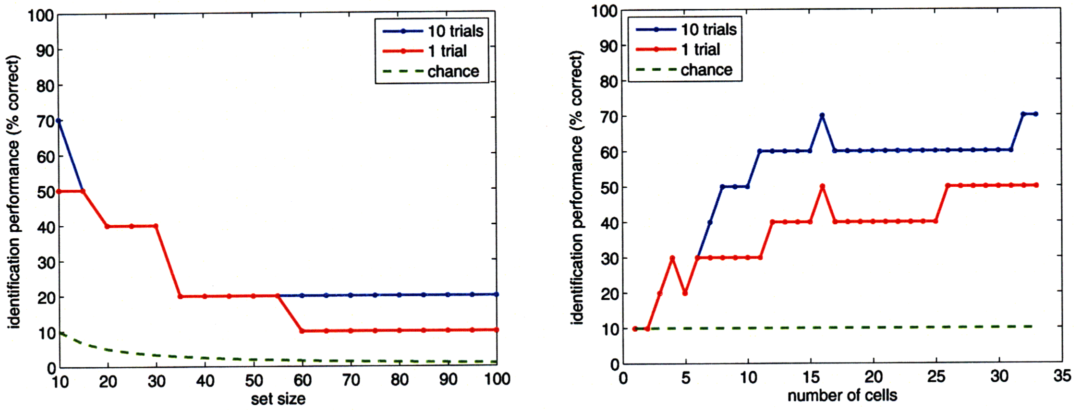

(a)

(b)

Figure 5-3: Factors that affect identification performance. Scaling of identification performance with *a*, set size, using all 33 available cells, and *b*, number of cells, using a set size of 10 images. The dashed green line indicates chance performance.

images, compared to our model trained over binary stimuli and using ML estimates from moving edges that confound space and time. Third, they used images spanning 20°×20° of the visual field, compared to our 8°×8° images (200 μm $\approx$ 1° for rabbit retina), and they used responses from 500 voxels compared to our set of 33 cells. We are therefore optimistic that our identification performance can be significantly improved.

5.3 Summary

In this chapter, we have applied the statistical inference algorithm presented in Chapter 3 to the problem of identifying a stationary natural image from a set of candidate images and have shown that the algorithm achieves moderate accuracy that scales with set size and number of cells.

Chapter 6

Discussion

In this thesis, we have presented an algorithm that is able to interpret ganglion cell firing patterns available to the brain to draw conclusions about the visual input. The linear-nonlinear Poisson neural model provides a simple intuitive way of determining the expected firing rate to arbitrary stimuli, and a cost function, here likelihood, takes into account the stochastic variability of spike responses. Both model and visual stimulus parameters can be found through optimization, and we have shown that even when using moving edge stimuli that intrinsically link spatial and temporal features, the estimated model parameters agree with the traditional STA approach and fall within physiologically reasonable ranges. In decoding the visual stimulus for the case of moving edges and natural scenes, moderate accuracy is achievable using the responses of a few cells. We achieved improved estimation of global motion parameters compared to previous algorithms, suggesting that the brain is able to encode visual parameters with a higher degree of fidelity than previously determined. Furthermore, we were able to identify a novel set of natural images with high accuracy despite ignoring LEDs and using a model trained only on binary stimuli.

6.1 Limitations

Despite the advantages of our approach, it fails to account for several known properties of retinal ganglion cell behavior. First, we have modeled the cellular system as a simple linear filter with point nonlinearity, but neurons fire based on complex mechanisms of propagating electrical activity. As such, cells do not obey superposition; instead they have different response characteristics when presented with, for example, a small box versus a large bar across their receptive fields. Cells are also not time-invariant, which is particularly evident when the stimulus remains constant over a long duration or consists of singularities. In the first case, the cells may adapt, possibly to conserve energy and not relay information about non-novel stimuli. In the second case, the exact analytic expressions using to determine the rate function become important because as Rodieck[23] explains, the convolution of the spatial and temporal functions in the model has a smoothing effect so that the exact expressions for $f(x, y)$ and $h(t)$ are minor compared to their general shapes. For stimuli of short duration, this smoothing is no longer in effect so that we can no longer look at the average response as a measure of expected response. Finally, the model does not account for physiological behavior at other retinal levels that may affect ganglion cells; for example, strong white light can bleach photoreceptors and distort RGC responses.

For the case of maximum likelihood, the Poisson assumption cannot be an accurate model of neural behavior since neural firing depends on past spiking history whereas a Poisson process is memoryless. A Poisson process cannot account for refractory periods from depolarization, nor does it support the observation that a neuron will often fire in doublets or triplets, as seen from its autocorrelation function in response to white noise. Experimental data has also shown that RGC ISIs have a longer tail than suggested by a Poisson distribution. On the other hand, because RGC spike trains are history dependent, the Poisson assumption is quite conservative and exhibits more variability across trials than experimentally observed data. [14] showed that when large numbers of trials are combined, the collective set of spike responses follows a Poisson process, but the spike response from an individual trial is non-Poisson. As our analysis in the decoding problem focuses on individual responses, we cannot average out noise across trials, not take advantage of the Poisson behavior that arises in collective responses across repeated trials.

Finally, in optimization, we attempt to account for all response variability through the cost function, whether this variability arises from the stochastic behavior of cells, noise introduced by the experimental measurements, or because it is compensating for errors of the model.

Again these assumptions were made for analytic tractability. We will discuss methods of overcoming some of these limitations in the next section, along with possible generalizations and avenues for future work.

6.2 Future Work

An obvious way of extending our method is to allow for a more generalized neural model or spiking process. For example, to incorporate feedback, we could use the integrate-and-fire model of [20]) or the inhomogeneous Markov interval process of [14], both of which have a firing rate that is dependent on the current time and the time since the last spike. The latter is relatively simple to incorporate as [3] showed how the time-rescaling theorem can be used to convert non-Poisson processes to homogeneous Poisson processes and then used to assess the goodness-of-fit of different models.

We can also take advantage of the observation that a spike that occurs in a cluster is more likely than a lone spike to be caused by an external stimulus. This reduces the effect of spontaneous background spikes and would result in a more robust optimization procedure. One problem with our current approach is that ignoring the spikes with the lowest likelihood could shift our focus in the wrong direction, particularly if the initial estimates are far from their true values. Preliminary studies using a likelihood function with weights dependent on interspike intervals have resulted in slightly more accurate estimates during visual stimulus decoding.

The neural model can also be generalized to allow characterization of other RGC types, for example directionally-selective cells or local edge detectors, thereby accounting for orientation or spatial frequency tuning. In a similar manner, it can be extended to characterize other sensory neurons, for example, using Gabor wavelets rather than a difference of Gaussians to account for the spatial responses of simple V1 cells. At the front-end, if we wish to incorporate more complicated, nonlinear visual stimulus dependence, we can pre-process the visual input to account for photoreceptor bleaching or use object representations rather than simple intensity values.

Synchrony and dependencies across neurons can be modeled by introducing networks to interconnect neural responses. For example, a spike in one neuron can promote or inhibit firing in a neighboring neuron. To keep the problem computationally tractable, as model parameters would have to be optimized collectively and expected responses would no longer be a direct function of the visual stimulus, we would have to restrict ourselves to simple networks, e.g. looking at relationships among pairs of cells.

Since the model is capable of determining the response to arbitrary stimuli, we are no longer restricted to simple artificial stimuli such as moving gratings or white noise to characterize cells. An interesting case would be to map neural behavior using complex classes of stimuli, such as natural scenes as this set most closely approximates the natural environment of the eye.

Finally, we have focused in this thesis on two specific decoding problems: estimating the global motion parameters of a moving edge, and identifying a natural scene. Is there a method that can *reconstruct* arbitrary stimuli? A solution proposed by [25] for cells in the cat lateral geniculate nucleus used linear filters, and there is agreement in the retinal community that RGC response to visual luminance is well-captured by a linear relationship[30, 25, 32]. It is possible to extend the model-based statistical inference framework of this thesis to reconstruct a scene in a computationally feasible manner by reducing the model to its associated linear operator. The neural response is then a realization of the operator output corrupted by Poisson noise. A similar problem has been studied in the context of image denoising when a charge-coupled device (CCD) is used to capture the image[17, 2], though these papers make the additional assumption that the number of measurements is equal to the image size. In both the neural and CCD perspective, the ML estimate of an arbitrary stimulus given a set of responses is an ill-posed problem. As before, one possibility is to restrict ourselves to natural scenes so as to take advantage of the inherent structure of such stimuli through regularization: for example, maximize the likelihood of the observed neural responses subject to the estimate achieving a low gradient. Preliminary results with stationary images suggest that this approach is capable of finding a reconstruction with lower minimum mean-square error than that produced by applying simple linear filter techniques.

6.3 Conclusion

The big question remains: how does the visual encoding system work at the layer of the retinal ganglion cells, and how can we accurately encode visual stimuli and decode spike responses? While we have shown that the LNP neural model can capture RGC behavior, at least for the cases of moving edges and natural scenes, we have also presented reasons why this model is woefully inadequate. Approaches that attempt to improve on the LNP model by increasing the parameter space can quickly become computationally intractable, and we suggest future work should focus on developing more physiological neural models, such as the integrate-and-fire model, that can capture the dynamics of RGC behavior without increasing computational complexity.

Appendix A

Experimental Procedure

This chapter outlines the experiments that were conducted to collect the data in chapters 4 and 5. This includes rabbit retinal tissue preparation, MEA setup, spike sorting, and designing the visual stimuli protocol. Parts of this chapter were taken with permission from [27, 26].

Experiments were conducted in collaboration with neurophysiologist Steven Stasheff in the Department of Pediatric Neurology and Neuro-opthalmology at the University of Iowa and with neurophysiologist Ofer Ziv in the Research Laboratory for Electronics at the Massachusetts Institute of Technology. In the results sections of this work, retinal pieces A, C, and F are from Stasheff's setup, and retinal pieces B, D, and E are from Ziv's setup.

A.1 Preparation

Whole-mounted pieces of rabbit retina were prepared and maintained in vitro for recording. Briefly, New Zealand white rabbits of either sex (3-5 kg) were anesthetized with xylazine (5-10 mg/kg) and ketamine (30-100 mg/kg) to the point that the corneal reflex was abolished. The animal was enucleated, the globe hemisected, and the vitreous removed. The animal was killed with an overdose of ketamine, according to a protocol approved by the respective university. Under infrared illumination to minimize exposure to visible light, using a dissecting microscope (Leica Microsystems, Bannockburn, IL) with infrared image intensifiers (BE Myers, Redman, WA), the retina was dissected from the retinal pigmentary epithelium and placed ganglion cell layer down onto a multielectrode recording array (8 × 8 planar grid with no corner electrodes, 10 µm contacts spaced 200 µm apart; Multichannel Systems, Reutlingen, Germany). This was maintained in a recording chamber attached to a microscope stage and superfused at 2.5-3.5 mL/min with oxygenated Ames' medium at 33-37.

A.2 Multielectrode Recording

A 60-channel amplifier (Multi-channel Systems, Reutlingen, Germany) was mounted on a microscope stage (Zeiss Axioplan, Göttingen, Germany) and interfaced with digital sampling hardware and software (Bionic Technologies, Inc., Salt Lake City, UT) for recording and analyzing spike trains from each of the electrodes in the array. Digitized data initially were streamed onto the computer's hard drive and further analyzed online. After transfer of the retina to the recording chamber, recordings were allowed to stabilize for $\geq$1 h as evidenced by stable action potential amplitudes, number of cells recorded, frequency of spontaneous firing, and consistency of light-evoked responses (where obtainable).

For the Stasheff setup, the MEA recording system samples waveforms at 30 kHz. If a digitized waveform exceeds a user-defined threshold, it is stored in memory along with its occurrence time. These thresholds (one for each channel/electrode) are set in such a way so as to minimize the recording of events other than action potentials. In this manner, only action potentials and their corresponding occurrence times are stored in memory; faulty waveforms are discarded. For the Ziv setup, the MEA recording system stores the digitized waveforms sampled at 25 kHz.

A.3 Spike Waveform Analysis

For the Stasheff setup, action potential (spike) waveforms accepted for further analysis were $\geq$60 µV in amplitude and $>$1.85 times the RMS of the background signal.

For the Ziv setup, the spike-sorting software Plexon (Plexon, Inc., Dallas, TX) was used to extract the waveforms ≥4 times the RMS of the background signal. In both setups, to distinguish responses from different cells that might appear on the same electrode, Plexon was used for supervised automated sorting of action potential profiles according to a principle components analysis (PCA) paradigm. For each electrode, the software displays a random sample of ≤10,000 spike waveforms along with all the two-dimensional projections of each waveform in the space defined by the first three principle components (eigenvectors computed from the correlation matrix for this data subset).

The individual waveforms were partitioned iteratively into 1-5 clusters according to an automated T-distribution expectation-maximization paradigm. This was followed by further manual assignment of waveforms to specific clusters to reduce automatic sorting errors. Appropriate assignment of individual waveforms to distinct cells was confirmed further by analysis of the corresponding spike trains. Interspike interval (ISI) histograms were computed for each spike train by measuring the intervals between spikes in the train for all possible spike pairs within a candidate cluster, then distributing these values in bins of 0.2 ms width. ISI histograms from accepted data demonstrated a refractory period of >1 ms(typically 2-5 ms) and did not reflect any of several patterns of recognizable noise: 60 Hz, very high frequency (>10 kHz) transients, or waveforms distinct from those of extracellular action potentials (e.g. sinusoidal oscillations).

A.4 Visual Stimuli

A miniature computer monitor (Lucivid, MicroBrightField, Colchester, VT) projected visual stimuli through a ×5 objective, and these were focused via standard microscope optics (Zeiss Axioplan) onto the photoreceptor layer of the retina. Luminance was calibrated via commercial software (VisionWorks, Vision Research Graphics, Durham, NH), using a photometer (Minolta, Ramsey, NJ) and photodiode (Hamamatsu S1133-11, Hamamatsu City, Japan) placed in the tissue plane. The refresh rate of the monitor (66 Hz) was chosen to avoid entrainment of retinal ganglion cells that might contaminate light responses. The same software controlled and recorded stimulus parameters, passing synchronization pulses to the data acquisition computer via a parallel interface with approximately 10 μs precision.

The retinal piece was stimulated with various stimuli which were crafted using commercially available software (Stasheff setup: VisionWorks, Vision Research Graphics, Durham, NH; Ziv setup: Adobe Premiere, Adobe Systems Inc., San Jose, CA)). The projected images are pixelated with 800 × 600 resolution on an area of 2286 μm × 1714 μm using the conversion 35 pixels = 100 μm (Stasheff setup), or pixelated with 640 × 480 resolution on an area of 2174 μm × 1680 μm (Ziv setup).

A.4.1 M-Sequence

A 16×16 square lattice of randomly flickering elements was presented for 18-25 min. Random flicker was created by varying the intensity of the display at each element according to an m-sequence[22]. The m-sequence consisted of 32,768 frames presented at 22 Hz (Stasheff setup) or 30 Hz (Ziv setup), with an element size of 100 μm.

A.4.2 Moving Edges

Moving bars of width equal to the screen width and contrast (±25%) were presented at a constant speed and direction of motion. The width of the bar was chosen to be sufficiently large so that a cell only sees the entering edge. The speeds and directions probed were (300 μm/s, 600 μm/s, 900 μm/s, 1200 μm/s) and (0°, ±90°, 180°), where the speeds were chosen based on speed sensitivity in the rabbit retina[11, 10, 31]. A run consisted of the following steps:

1. The screen is fully dark for 4 s.

2. A bright bar enters the screen from the left and moves to the right at a speed of 300 μm/s.

3. Once the screen has become fully bright, it stays fully bright for 4 s. This is the (ON, 300 µm/s, 0°) stimulus.
4. The bright bar recedes from the screen at a speed of 300 µm/s. This can also be viewed as a dark bar entering the screen from the right and moving to the left, hence this is the (OFF, 300 µm/s, 180°) stimulus.
5. The run ends when the screen has become fully dark.

The above is repeated at the same speed with the bright bar moving along the other three directions (entering from the right and receding, entering from the top and receding, entering from the bottom and receding). A sequence at a single speed therefore consisted of the following stimuli, separated by 4s delays where the screen is fully dark or fully bright:

1. (ON, 300 µm/s, 0°)
2. (OFF, 300 µm/s, 180°)
3. (ON, 300 µm/s, 180°)
4. (OFF, 300 µm/s, 0°)
5. (ON, 300 µm/s, −90°)
6. (OFF, 300 µm/s, 90°)
7. (ON, 300 µm/s, 90°)
8. (OFF, 300 µm/s, −90°)

This sequence is then repeated using the other three directions (600 µm/s, then 900 µm/s, then 1200 µm/s) to form a trial consisting of all speed and direction pairs. Finally, a single experiment consisted of 10-30 trials.

Also note that the rasterization of the CRT display introduced a space-time sampled approximation of a moving bar. For example, for the Stasheff setup, a bar nominally moving at 1200 µm/s was in fact redrawn on the CRT every 15.15 ms displaced by 18.18 µm. As noted in [7], the effect of this discretization was probably small.

A.4.3 Natural Images

The protocol for natural images closely follows that of [15]. Photographs from the Berkeley Segmentation Dataset (`http://www.eecs.berkeley.edu/Research/`

Projects/CS/vision/grouping/segbench) were converted to grayscale, cropped to the central 576×576 pixels, and linearly transformed so that the 1/10th and 99 9/10th percentiles of the original pixel values mapped to the minimum (0) and maximum (255) pixel values. This was downsampled to 96×96 pixels by block averaging over non-overlapping 6×6 squares, and the resulting image centered on a 800×600 pixel gray (128) background. Each image pixel (not machine pixel) therefore covered an area of approximately 16.3 μm×16.3 μm (Ofer setup).

The images were presented in successive 4 s trials, with the order determined by the segmentation dataset. A trial consisted of an image being presented for 1 s followed by presentation of a gray background for 3 s. Each 1 s image presentation consisted of the image being flashed ON-OFF-ON-OFF-ON where ON corresponds to presentation of the image for 200 ms, and OFF corresponds to presentation of the gray background for 200 ms.

Appendix B

Calculation of Rate Functions

Here we derive the rate function for a moving edge stimulus, as seen in Section 3.2.5. As a reminder, the rate function is given by Equation (3.2), repeated below,

$$\lambda(t) = n\left(\int_{-\infty}^{+\infty}\int_{-\infty}^{+\infty} f(x,y)s(x,y,t)dxdy * h(t)\right), \tag{B.1}$$

and our chosen neural functions are Equations (3.3), (3.5),and (3.7), also repeated below.

$$f(\underline{x}) = \pm\left(\frac{1}{2\pi|\mathbf{\Sigma}|^{1/2}}\, e^{-\frac{1}{2}\left(\underline{x}-\underline{\mu}\right)^{T}\mathbf{\Sigma}^{-1}\left(\underline{x}-\underline{\mu}\right)} - k\frac{1}{2\pi|r\mathbf{\Sigma}r|^{1/2}}\, e^{-\frac{1}{2}\left(\underline{x}-\underline{\mu}\right)^{T}(r\mathbf{\Sigma}r)^{-1}\left(\underline{x}-\underline{\mu}\right)}\right) \tag{B.2}$$

$$h_s(t) = \alpha^2\left[(t-\tau)e^{-\alpha(t-\tau)} + m\right]u(t-\tau) \tag{B.3}$$

$$n(g) = \max(ag, d) \tag{B.4}$$

For simplicity, we will focus solely on the excitatory component of the spatial sensitivity function

$$f(x,y) = \pm N(x,y;\underline{\mu},\mathbf{\Sigma}), \tag{B.5}$$

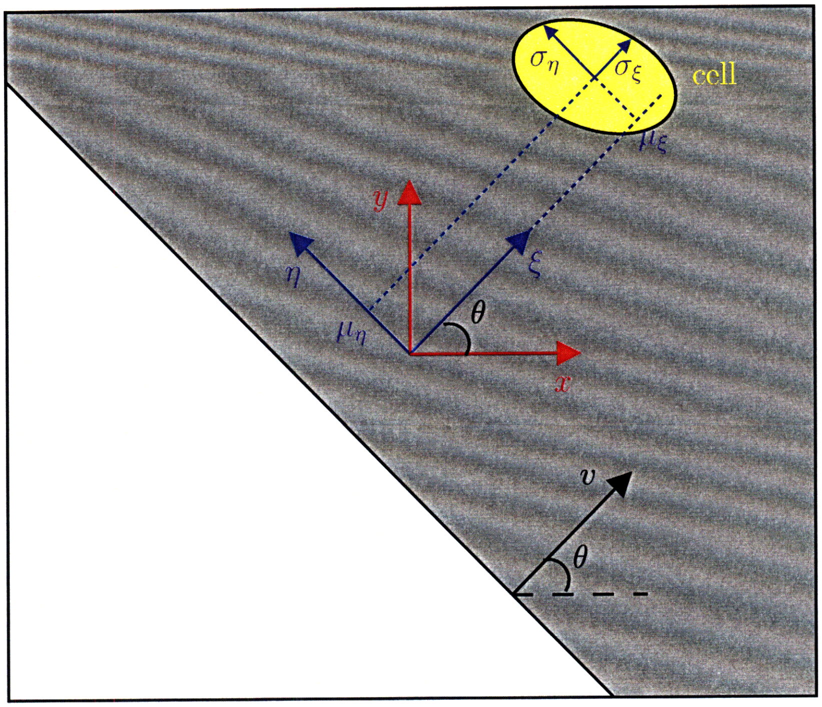

Figure B-1: Moving edge stimulus. A bar of constant intensity moves at a constant speed v and in a constant direction θ. The coordinate transform and spatial parameters of the cell are shown.

and generalize our results at the end of this chapter. We will also use the following intermediate signals in analyzing our rate function:

$$l(t) = \int_{-\infty}^{+\infty} \int_{-\infty}^{+\infty} f(x,y)s(x,y,t)dxdy \tag{B.6}$$

$$g(t) = l(t) * h(t) \tag{B.7}$$

Since we have defined the temporal sensitivity function in terms of the step response, we will use the alternative formulation

$$\dot{l}(t) = \int_{-\infty}^{+\infty} \int_{-\infty}^{+\infty} f(x,y)\dot{s}(x,y,t)dxdy \tag{B.8}$$

$$g(t) = \dot{l}(t) * h_s(t). \tag{B.9}$$

For the moving edge stimulus, let $\dot{s}(x,y,t)$ represent the gradient of the stimulus intensity, and use the rotated coordinate system (ξ, η), where ξ and η are orthogonal but ξ is aligned to the direction of motion of a moving edge (see Figure B-1). Then

$$\dot{s}(\xi, \eta, t) = \pm\, \delta(t - \xi/v). \tag{B.10}$$

Note that the stimulus is now independent of η. For simplicity, we focus on the response of an ON cell to an ON stimulus so that we can drop the ± notation in the spatial sensitivity function and the stimulus.

Using the Step Response

Before continuing, we should prove that the our representation of the stimulus and the temporal sensitivity function in terms of the step response is equivalent to the standard representation in terms of the impulse response. To do this, let $\dot{x}(t)$ represent the derivative of the signal $x(t)$ with respect to t. Then

$$\begin{aligned}\int_{-\infty}^{\infty} \tfrac{d}{dt}\left(a(\tau)b(t-\tau)\right) d\tau &= \int_{-\infty}^{\infty} \dot{a}(\tau)b(t-\tau)d\tau - \int_{-\infty}^{\infty} a(\tau)\dot{b}(t-\tau)d\tau \\ &= \dot{a}(t) * b(t) - a(t) * \dot{b}(t),\end{aligned}$$

and in the limit,

$$\int_{-\infty}^{\infty} \tfrac{d}{dt}\left(a(\tau)b(t-\tau)\right) d\tau = \lim_{t\to\infty} a(t)b(-t) - a(-t)b(t)$$

Thus, provided that $\lim_{t\to\infty} a(t)b(-t) - a(-t)b(t) = 0$, we have

$$\dot{a}(t) * b(t) = a(t) * \dot{b}(t).$$

To tie this back into the neural model, let $a(t) = l(t)$ and $b(t) = h_s(t)$ so that

$$l(t) * h(t) = \dot{l}(t) * h_s(t).$$

That is, given a signal $l(t)$ and a system impulse response $h(t)$, the convolution of the signal with the system impulse response is equivalent to convolution of the gradient of the signal with the system step response. We can verify that the our limiting condition holds since by design, we have $l(t) \to 0$ for $t \to \pm\infty$ and $h_s(t) \to 0$ for $t \to -\infty$. Furthermore, if $m = 0$, $h_s(t) \to 0$ for $t \to \infty$.

B.1 Effect of the Spatial Sensitivity Function

Our first step is to determine the effect of the spatial sensitivity function, so we focus our efforts on finding $\dot{l}(t)$. To simplify the mathematics, we use the trick mentioned in Section 3.2.1. That is, we rotate $f(x, y)$ to the new coordinate system (ξ, η) defined above to give the rotated spatial function

$$f'(\xi,\eta) = N(\xi,\eta;\underline{\mu}',\Sigma'). \tag{B.11}$$

To calculate the new mean and standard deviation, we use the rotation matrix

$$\mathbf{R} = \begin{bmatrix} \cos(\theta) & \sin(\theta) \\ -\sin(\theta) & \cos(\theta) \end{bmatrix}, \tag{B.12}$$

so that

$$\underline{\mu}' = \mathbf{R}\underline{\mu} = [\mu_\xi, \mu_\eta]^T \tag{B.13}$$

$$\Sigma' = \mathbf{R}\Sigma\mathbf{R}^T = \begin{bmatrix} \sigma_\xi^2 & \sigma_{\xi\eta} \\ \sigma_{\xi\eta} & \sigma_\eta^2 \end{bmatrix} \tag{B.14}$$

where

$$\mu_\xi = \mu_x \cos(\theta) + \mu_y \sin(\theta) \tag{B.15}$$

$$\mu_\eta = -\mu_x \sin(\theta) + \mu_y \cos(\theta) \tag{B.16}$$

$$\sigma_\xi^2 = \sigma_x^2 \cos^2(\theta) + 2\sigma_{xy} \cos(\theta)\sin(\theta) + \sigma_y^2 \sin^2(\theta) \tag{B.17}$$

$$\sigma_{\xi\eta} = -\sigma_x^2 \cos(\theta)\sin(\theta) + \sigma_{xy}(\cos^2(\theta) - \sin^2(\theta)) + \sigma_y^2 \cos(\theta)\sin(\theta) \tag{B.18}$$

$$\sigma_\eta^2 = \sigma_x^2 \sin^2(\theta) - 2\sigma_{xy} \cos(\theta)\sin(\theta) + \sigma_y^2 \cos^2(\theta). \tag{B.19}$$

Since $s(\xi, \eta, t)$ is independent of η, rather than integrating $f(x, y)$ jointly over the dependent limits of x and y, we can integrate $f'(\xi, \eta)$ over the independent ranges of ξ and η. Then

$$\begin{aligned}
\dot{l}(t) &= \int_{-\infty}^{+\infty} \int_{-\infty}^{+\infty} f(x, y)\, \dot{s}(x, y, t)\, dx dy \\
&\overset{(a)}{=} \int_{-\infty}^{+\infty} \int_{-\infty}^{+\infty} f'(\xi, \eta)\, \dot{s}'(\xi, \eta, t)\, d\xi d\eta \\
&\overset{(b)}{=} \int_{-\infty}^{+\infty} \int_{-\infty}^{+\infty} N(\xi, \eta; \mu', \Sigma')\, \delta(t - \xi/v)\, d\xi d\eta \\
&\overset{(c)}{=} \int_{-\infty}^{+\infty} N(\xi; \mu_\xi, \sigma_\xi)\, \delta(t - \xi/v)\, d\xi \\
&\overset{(d)}{=} |v| N(vt; \mu_\xi, \sigma_\xi),
\end{aligned} \tag{B.20}$$

where (a) changes the coordinate system, (b) substitutes Equations (B.11) and (B.10), (c) marginalizes the Gaussian over η, and (d) takes the line integral of the marginalized Gaussian. To simplify notation, we will drop the ξ subscripts in subsequent equations, with the understanding that μ and σ represent the mean and standard deviation of the Gaussian in the ξ-direction.

B.2 Effect of the Temporal Sensitivity Function

Now we determine the effect of the temporal sensitivity function, focusing our efforts on finding $g(t)$. For $t \geq \tau$,

$$
\begin{aligned}
g(t) &= \dot{l}(t) * h_s(t) \\
&\overset{(a)}{=} |v| N(vt; \mu, \sigma) * \alpha^2 \left[(t-\tau)e^{-\alpha(t-\tau)} + m\right] u(t-\tau) \\
&\overset{(b)}{=} \int_{-\infty}^{t-\tau} |v| N(vt'; \mu_\xi, \sigma_\xi) \cdot \alpha^2 \left[(t-t'-\tau)e^{-\alpha(t-t'-\tau)} + m\right] dt' \\
&\overset{(c)}{=} \int_{-\infty}^{v(t-\tau)} N(x; \mu, \sigma) \cdot \alpha^2 \left[(t-x/v-\tau)e^{-\alpha(t-x/v-\tau)} + m\right] dx \\
&\overset{(d)}{=} \int_{-\infty}^{x'} N(x; \mu, \sigma) \cdot \alpha^2 \left[\left(\tfrac{x'-x}{v}\right) e^{-\alpha\left(\frac{x'-x}{v}\right)} + m\right] dx \\
&= \underbrace{\tfrac{\alpha^2 x'}{v} \int_{-\infty}^{x'} e^{-\alpha\left(\frac{x'-x}{v}\right)} N(x; \mu, \sigma)\, dx}_{(1)} - \underbrace{\tfrac{\alpha^2}{v} \int_{-\infty}^{x'} x\, e^{-\alpha\left(\frac{x'-x}{v}\right)} N(x; \mu, \sigma)\, dx}_{(2)} \\
&\qquad \underbrace{+\, \alpha^2\, m \int_{-\infty}^{x'} N(x; \mu, \sigma)\, dx,}_{(3)}
\end{aligned}
$$

where (a) substitutes Equations (B.20) and (B.3), (b) uses the convolution formula, (c) performs the change of variables $x = vt'$, and (d) substitutes $x' = v(t-\tau)$. After some algebra,

$$
\begin{aligned}
(1) &= \tfrac{\alpha^2 x'}{v}\, e^{-\beta}\, \Phi(x'; \mu + z, \sigma) \\
(2) &= \tfrac{\alpha^2}{v}\, e^{-\beta} \left[-\sigma^2 N(x'; \mu + z, \sigma) + (\mu + z)\Phi(x'; \mu + z, \sigma)\right] \\
(3) &= m\, \alpha^2\, \Phi(x'; \mu, \sigma),
\end{aligned}
$$

where $z = \sigma^2\alpha/v$ and $\beta = \frac{\alpha}{v}\left[x' - (\mu + z/2)\right]$ come from completing the square in the exponential. Pulling these three components together,

$$g(t) = (g_m + g_t)\, u(t-\tau), \tag{B.21}$$

where

$$
\begin{aligned}
g_m &= (3) = m\,\alpha^2\,\Phi(x';\mu,\sigma)\\
g_t &= (1)+(2) = \tfrac{\alpha^2}{v}\,e^{-\beta}\,\left[(x'-(\mu+z))\,\Phi(x';\mu+z,\sigma)+\sigma^2\,N(x';\mu+z,\sigma)\right]
\end{aligned}
$$

represent the maintained and transient components of the response, respectively.

B.3 Analysis

For most generator values, $ax > d$ so that our nonlinearity function (B.4) can be thought of as a linear multiplier. Thus, we restrict our analysis here to how the different components of $g(t)$ behave as t and v increase.

B.3.1 Maintained versus Transient Response

As $t \to \infty$ (or equivalently, as $x' \to \infty$), we expect the maintained response to remain, and the transient response to die away. This is easily seen to be the case for the maintained component since $\lim_{x'\to\infty}\Phi(x';\mu,\sigma) = 1$ so that $\lim_{x'\to\infty} g_m = m\alpha^2$. The transient component is a bit harder to analyze. However, noting that $\exp(-\beta)$ dominates the transient behavior for large x', we see that $\lim_{x'\to\infty}\beta = \infty$ so that $\lim_{x'\to\infty} g_t = \lim_{x'\to\infty}\exp(-\beta) = 0$, and again, our expectations are satisfied.

B.3.2 Spatial versus Temporal Sensitivity

For large v (high velocity edges), the stimulus sweeps over the RF very quickly, so we expect that the chosen shape of the spatial sensitivity function should not be very influential. That is, we expect $l(t)$ to approach an impulse and $g(t)$ to take the shape of the temporal sensitivity function. Similarly, for small v (low velocity edges), the stimulus sweeps over the RF very slowly, so we expect that the chosen shape of the spatial sensitivity function should be very influential, and $l(t)$ and $g(t)$ should approach the shape of the spatial sensitivity function.

To see if this is true, we analyze the maintained and transient components separately. For high velocity edges, we have

$$\lim_{v \to \infty} g_m = \lim_{v \to \infty} m\, \alpha^2\, \Phi(v(t-\tau); \mu, \sigma) = m\alpha^2,$$

and if we rewrite g_t as

$$\begin{aligned} g_t &= \tfrac{\alpha^2}{v}\, \exp\left(-\tfrac{\alpha}{v}\left[v(t-\tau) - \left(\mu + \tfrac{\sigma^2\alpha}{2v}\right)\right]\right) \times \\ &\quad \times \left[\left(v(t-\tau) - \left(\mu + \tfrac{\sigma^2\alpha}{v}\right)\right) \Phi\left(x'; \mu + \tfrac{\sigma^2\alpha}{v}, \sigma\right) + \sigma^2\, N\left(v(t-\tau); \mu + \tfrac{\sigma^2\alpha}{v}, \sigma\right)\right] \\ &= \alpha^2\, \exp\left(-\alpha\left[(t-\tau) - \left(\tfrac{\mu}{v} + \tfrac{\sigma^2\alpha}{2v^2}\right)\right]\right) \times \\ &\quad \times \left[\left((t-\tau) - \left(\tfrac{\mu}{v} + \tfrac{\sigma^2\alpha}{v^2}\right)\right) \Phi\left(x'; \mu + \tfrac{\sigma^2\alpha}{v}, \sigma\right) + \tfrac{\sigma^2}{v}\, N\left(v(t-\tau); \mu + \tfrac{\sigma^2\alpha}{v}, \sigma\right)\right], \end{aligned}$$

then

$$\lim_{v \to \infty} g_t = \alpha^2 \exp\left(-\alpha(t-\tau)\right)(t-\tau).$$

Thus,

$$\begin{aligned} \lim_{v \to \infty} g(t) &= \lim_{v \to \infty} (g_m + g_t)\, u(t-\tau) \\ &= \alpha^2\left[(t-\tau)e^{-\alpha(t-\tau)} + m\right] u(t-\tau) \\ &= h_s(t), \end{aligned}$$

as we hypothesized. Analysis for low velocity edges becomes quite complicated, but we can similarly prove to ourselves that $g(t)$ approaches $\dot{l}(t)$.

B.4 Sample Output Functions

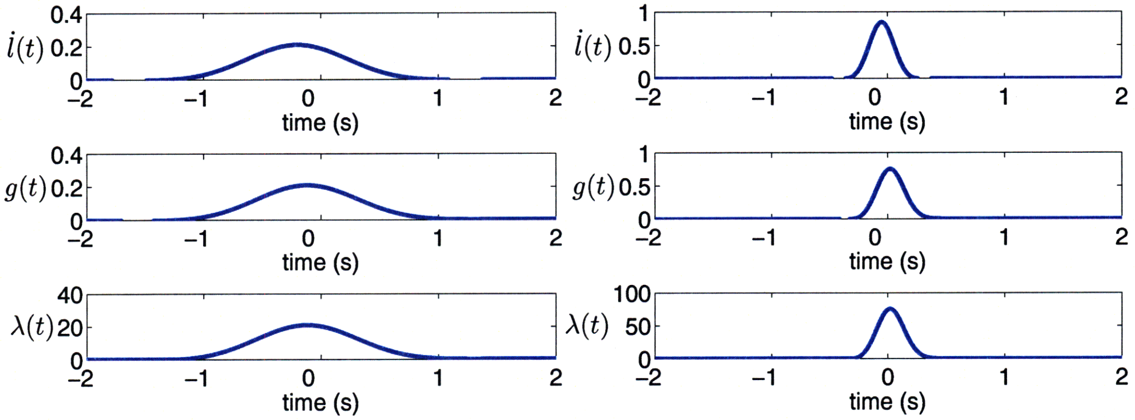

(a) (ON,300µm/s, 0°) (b) (ON,1200µm/s, 0°)

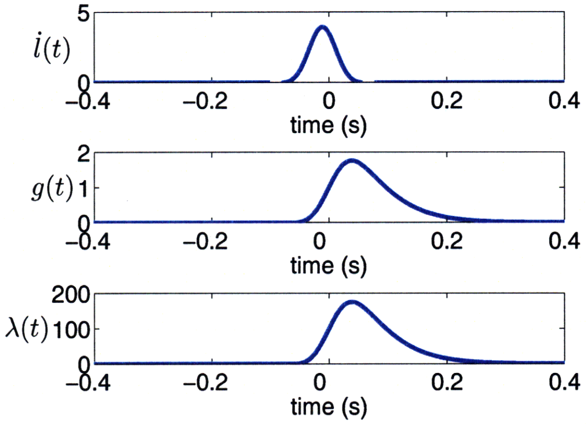

(c) (ON,5600µm/s, 0°)

Figure B-2: Rate function outputs. The neural model used the functions (B.2)-(B.4), with parameters optimized through ML. The responses are for the ON cell corresponding to electrode 16, unit 2 in retinal piece B. (See Section 4.1 for the procedure used to determine the optimal neural parameters, and Section 4.2 for more detail on the associated experiment.) The visual stimuli were moving edges with the shown (polarity, v, θ) parameters. Note the differing time scales and how the generator signal $g(t)$ changes for increasing v. In particular, it becomes narrower in time, higher in peak amplitude, and takes on a longer tail. Also, for small v, $g(t)$ is similar to the Gaussian of $f(x, y, t)$, and for large v, $g(t)$ is similar to the decaying exponential of $h_s(t)$.

B.5 Generalizing the Neural Model Functions and the Stimulus

In general, the double integral and convolution of (B.6) cannot be solved in closed form. Closed form solutions were possible in our derivations in Sections B.1 and B.2 since the standard normal spatial function and decaying exponential temporal function, as well as their product, are all integrable. This means that in general, we can pick neural sensitivity functions

$$f(x,y) = \sum_k g_k N(x,y;\underline{\mu}_k,\Sigma_k) \tag{B.22}$$

$$h_s(t) = \sum_k (a_k + b_k t)e^{-(c_k t + d_k t^2)}u(t-t_k) \tag{B.23}$$

and stimuli with gradient $\dot{s}(x,y,t)$ that satisfy

$$\int_{-\infty}^{\infty}\int_{-\infty}^{\infty} f(x,y,t)\dot{s}(x,y,t)dxdy = \sum_k (a'_k + b'_k t)e^{-(c'_k t + d'_k t^2)+m_k}u(t-t_k) \tag{B.24}$$

and still arrive at closed form solutions. In particular, this allows us to use a difference of Gaussians spatial response, and it allows for more complicated visual stimuli such the moving bars with a Gaussian profile used by [7], or more general sinusoidal gratings or checkerboard patterns.

Bibliography

[1] Riccardo Barbieri, Michael C. Quirk, Loren M. Frank, Matthew A. Wilson, and Emory N. Brown. Construction and analysis of non-poisson stimulus-response models of neural spiking activity. *Journal of Neuroscience Methods*, 105(1):25–37, January 30, 2001.

[2] Johnathan M. Bardsley. A theoretical framework for the regularization of poisson likelihood estimation problems. Technical Report 25, University of Montana, Department of Mathematical Sciences, 2008.

[3] Emory N. Brown, Riccardo Barbieri, Val e Ventura, Robert E. Kass, and Loren M. Frank. The time-rescaling theorem and its application to neural spike train data analysis. *Neural Computation*, 14(2):325–346, 2002.

[4] E. J. Chichilnisky. A simple white noise analysis of neuronal light responses. *Network: Computation in Neural Systems*, 12:199–213, 2001.

[5] E. J. Chichilnisky and R. S. Kalmar. Functional asymmetries in on and off ganglion cells of primate retina. *Journal of Neuroscience*, 22:2737–2747, 2002.

[6] Stephen H. Devries and Denis A. Baylor. Mosaic arrangement of ganglion cell receptive fields in rabbit retina. *Journal of Neurophysiology*, 78:2048–2060, 1997.

[7] E. S. Frechette, A. Sher, M. I. Grivich, D. Petrusca, A. M. Litke, and E. J. Chichilnisky. Fidelity of the ensemble code for visual motion in primate retina. *Journal of Neurophysiology*, 94(1):119–135, 2005.

[8] Stephen Grossberg, Ennio Mingolla, and William D. Ross. Visual brain and visual perception: how does the cortex do perceptual grouping? *Trends in Neurosciences*, 20(3):106–111, February 19, 1997.

[9] K. Shane Guillory, Shy Shoham, and Richard A. Normann. Discrete stimulus estimation from neural responses in the turtle retina. *Vision Research*, 46(12):1876–1885, June 2006.

[10] Shigang He, W. R. Levick, and D. I. Vaney. Distinguishing direction selectivity from orientation selectivity in the rabbit retina. *Visual Neuroscience*, 15(3):438–447, 1998.

[11] Shigang He and William R. Levick. Spatial-temporal response characteristics of the on-off direction selective ganglion cells in the rabbit retina. *Neuroscience Letters*, 285(1):25–28, 2000.

[12] David Hubel. *Eye, Brain, and Vision.* Number 22 in Scientific American Library. W. H. Freeman, May 15, 1995.

[13] Mehrdad Jazayeri and J. Anthony Movshon. Optimal representation of sensory information by neural populations. *Nature Neuroscience*, 9(5):690–696, May 2006.

[14] Robert E. Kass and Valérie Ventura. A spike-train probability model. *Neural Computation*, 13(8):1713–1720, August 2001.

[15] Kendrick N. Kay, Thomas Naselaris, Ryan J. Prenger, and Jack L. Gallant. Identifying natural images from human brain activity. *Nature*, 452:352–356, March 19, 2008.

[16] J. J. Kulikowski, S. Marčelja, and P. O. Bishop. Theory of spatial position and spatial frequency relations in the receptive fields of simple cells in the visual cortex. *Biological Cybernetics*, 43(3):187–198, April 1982.

[17] Triet Le, Rick Chartrand, and Thomas J. Asaki. A variational approach to reconstructing images corrupted by poisson noise. *Journal of Mathematical Imaging and Vision*, 27(3):257–263, April 2007.

[18] S. Marčelja. Mathematical description of the responses of simple cortical cells. *Journal of the Optical Society of America*, 70(11):1297–300, November 1980.

[19] Nancy E. Medeiros and Christine A. Curcio. Preservation of ganglion cell layer neurons in age-related macular degeneration. *Investigative Ophthalmology and Visual Science*, 42(3):795–803, March 2001.

[20] Jonathan W. Pillow. *Bayesian Brain: Probabilistic Approaches to Neural Coding*, chapter Likelihood-Based Approaches to Modeling the Neural Code, pages 53–70. MIT Press, 2007.

[21] Jonathan W. Pillow, Liam Paninski, Valerie J. Uzzell, Eero P. Simoncelli, and E. J. Chichilnisky. Prediction and decoding of retinal ganglion cell responses with a probabilistic spiking model. *Journal of Neuroscience*, 25(47):11003–11013, November 23, 2005.

[22] R. C. Reid, J. D. Victor, and R. M. Shapley. The use of m-sequences in the analysis of visual neurons: Linear receptive field properties. *Visual Neuroscience*, 14:1015–1027, 1997.

[23] R. W. Rodieck. Quantitative analysis of cat retinal ganglion cell response to visual stimuli. *Vision Research*, 5(12):583–601, December 1965.

[24] Robert E. Soodak. Two-dimensional modeling of visual receptive fields using gaussian subunits. *Proceedings of the National Academy of Sciences of the United States of America*, 83(23):9259–9263, December 1, 1986.

[25] Garrett B. Stanley, Fei F. Li, and Yang Dan. Reconstruction of natural scenes from ensemble responses in the lateral geniculate nucleus. *Journal of Neuroscience*, 19(18):8036–8042, September 15, 1999.

[26] Steven F. Stasheff. Emergence of sustained spontaneous hyperactivity and temporary preservation of off responses in ganglion cells of the retinal degeneration (rd1) mouse. *Journal of Neurophysiology*, 99:1408–1421, 2008.

[27] Steven F. Stasheff and Richard H. Masland. Functional inhibition in direction-selective retinal ganglion cells: Spatiotemporal extent and intralaminar interactions. *Journal of Neurophysiology*, 88:1026–1039, 2002.

[28] Stavros Valavanis. Algorithms for estimating visual motion parameters from ganglion cell responses. Master's thesis, MIT, Cambridge, MA, January 2008.

[29] Jonathan D. Victor. Spike train metrics. *Current Opinion in Neurobiology*, 15(5):585–592, October 2005.

[30] David K. Warland, Pamela Reinagel, and Markus Meister. Decoding visual information from a population of retinal ganglion cells. *Journal of Neurophysiology*, 78(5):2336–2350, November 1997.

[31] Harry J. Wyatt and Nigel W. Daw. Directionally sensitive ganglion cells in the rabbit retina: Specificity for stimulus direction, size, and speed. *Journal of Neurophysiology*, 38(3):613–626, 1975.

[32] Günther M. Zeck, Xiao Quan, and Richard H. Masland. The spatial filtering properties of local edge detectors and brisk-sustained retinal ganglion cells. *Journal of Neuroscience*, 22(8):2016–2026, 2005.