Available online at www.sciencedirect.com
SCIENCE DIRECT
Vision Research 46 (2006) 1876–1885
Vision Research
www.elsevier.com/locate/visres

Discrete stimulus estimation from neural responses in the turtle retina

K. Shane Guillory [a,*], Shy Shoham [b], Richard A. Normann [a]

[a] Center for Neural Interfaces, Department of Bioengineering, University of Utah, Salt Lake City, UT 84112, USA
[b] Department of Biomedical Engineering, the Technion-Israel Institute of Technology, Haifa, Israel 32000

Received 21 December 2002; received in revised form 14 March 2005

Communicated by Robert Barlow

Abstract

In this paper, we investigate the decoding of flashed, full-field visual stimuli while recording from a population of retinal ganglion cells. We present a direct statistical method for determining the likelihood that a response was evoked by a particular stimulus, and use this method to estimate stimuli based on microelectrode array recordings in the turtle retina. This method uses the well-known time-varying Poisson model of neural firing, along with extensions to accommodate neural refractory periods. Unlike other approaches commonly used for Poisson processes, the specific formulation presented here is bin free and requires few user-specified parameters. Statistical dependency issues and the effects of stationarity on the estimation method are also discussed.

Keywords: Stochastic; Retina; Model; Stimulus; Discrimination; Turtle (*Pseudemys scripta elegans*)

1. Introduction

The entire visual experience of vertebrates is conveyed in the spatiotemporal patterns of action potentials that are output from their retinal ganglion cells. Early research into retinal encoding determined that ganglion cells can be placed into broad response classes (Granda & Fulbrook, 1989; Hartline, 1938) and that they had relatively localized, structured receptive fields (Kuffler, 1953). However, the responses of individual ganglion cells can exhibit significant variability to identical sets of visual stimuli and similar responses for very different stimuli (Reich, Victor, Knight, Ozaki, & Kaplan, 1997). These responses can also include stochastic variations due to the inherent noise in phototransduction and neural transmission processes, especially at near threshold levels. Because of these variabilities and ambiguities, it is generally recognized that populations of ganglion cells are required to reliably encode the visual scene. Nevertheless, the specifics of how these cells work together and how their firing patterns can best be interpreted are still the subject of much investigation. The exploration of methods for estimating stimuli based on the neural response is one approach to understanding these processes.

Neural responses to discrete stimuli are commonly analyzed with Peri-Stimulus Time Histograms (PSTHs). These plots average out the per-trial variability and provide the prototypical responses for a particular stimulus. Based on observed differences in the responses, one can create a variety of bin counting and vector representations that allow the stimuli to be determined from the neural response (Awiszus, 1997; Becker & Kruger, 1996; Gawne & Richmond, 1993; Geisler & Albrecht, 1997; Oram, Foldiak, Perrett, & Sengpiel, 1998; Salinas & Abbott, 1994). Ideally, however, a method is desired for directly evaluating the likelihood that the spikes from a single trial came from a certain prototypical response without ad hoc vector representations (Sanger, 2002).

Statistical methods for point processes can provide insight and a mathematical framework for studying this class of problems (Brown, Barbieri, Eden, & Frank, 2003; Johnson, 1996; Kass & Ventura, 2001; Perkel, Gerstein,

* Corresponding author.
E-mail address: shane.guillory@m.cc.utah.edu (K.S. Guillory).

0042-6989/$ - see front matter © 2005 Elsevier Ltd. All rights reserved.
doi:10.1016/j.visres.2005.12.001

& Moore, 1967). They provide a language for describing both the distribution of possible responses to stimuli and the likelihoods of different stimuli given an observed neural firing pattern. In this paper, we present a method for estimating the most likely stimuli among discrete sets on the basis of the neural response, and we apply this method to the responses of turtle retina ganglion cells. For the forward encoding model of the time-varying spike response, this method uses the well-known inhomogeneous Poisson process combined with a refractory renewal period following each spike. This combined model has also been called the Inhomogeneous Markov Interval or IMI process by Kass and Ventura (2001). Unlike most Poisson estimation methods used for neural signal decoding (Rieke, Warland, de Ruyter van Steveninck, & Bialek, 1997), the presented method is essentially bin-free and can include neural refractory periods in the estimation. Using this approach, we show that a discrete set of eight color stimuli can be decoded with an accuracy of 78% with recordings from a population of 18 cells. Performances from simulated cell populations constructed from the data are also presented.

When stochastic estimation methods are applied to a neural system, the statistical dependencies between cells and the stationarity of the responses over time must be determined (Johnson, 1996). Several researchers have proposed the existence of higher-order or synchrony codes that could be used among groups of cells in a variety of cortical and sensory neural systems (Abeles, 1991; Softky, 1995), including the retina (Meister, Lagnado, & Baylor, 1995). In this study, we used the normalized Joint PSTH (Aertsen, Gerstein, Habib, & Palm, 1989) to examine the correlation structure of the responses. This analysis found no significant correlations, and the cells were therefore treated as statistically independent coders of information for the estimation analysis. However, generalizations of the estimation method for correlated spike trains are discussed. In this study, we also examine the impact of the stationarity of the recorded responses on the model building and estimation performance.

2. Methods

2.1. Preparation

Recordings were made from ganglion cells in isolated turtle (*pseudemys scripta elegans*) retinas with isolation performed as described by Perlman, Normann, Chandler, and Lipetz (1990). In these experiments, a 100-electrode extracellular array with 1.5 mm electrodes in a 10×10 grid with 400 μm inter-electrode spacing was used. The retina was placed on a glass slide (photoreceptor side down), held in place by a millipore filter border, and superfused with an oxygenated (95% O_2, 5% CO_2) buffer solution (110 mM NaCl, 2.6 mM KCl, 2.0 mM $CaCl_2$, 2.0 mM $MgCl_2$, 22 mM $NaHCO_3$, and 10 mM D-glucose) delivered at 0.5 ml/min. Light stimuli were provided by a Hitachi Superscan Pro 620 monitor with a vertical refresh rate of 100 Hz and stimulus updates were performed between refreshes. The monitor image was focused by a 55 mm camera lens (f2.8) and prism system to produce a 6×6 mm image on the photoreceptor layer of the retina. Once the retina was in place, the electrode array was lowered into the ganglion cell layer of the retina until single unit activity became apparent. The data acquisition system allowed simultaneous online extraction of the spike timing and waveforms from all 100 electrodes in the array (Guillory & Normann, 1999). Multi-unit recordings were obtained and the single unit waveforms were classified offline using MATLAB implementations of the clustering algorithm described by Shoham, Fellows, and Normann (2003).

While unit activity was recorded, full-field light stimuli were presented in trials consisting of a 200 ms ON period followed by a 300 ms OFF period before the next trial. The light stimuli were randomly selected from an equally probable discrete set of eight stimuli composed of 100:1 contrast, ON–OFF binary combinations of the red, green, and blue channels of the monitor (ON intensities of 1.4, 2.0, and 1.5 mW/m^2, respectively). These intensities and color variations were selected as a simple stimulus set that could be differentiated by the pentachromatic visual system of the turtle. These color combinations appear to the humans as black (no stimulus), red, green, blue, cyan, magenta, yellow, and white. Data were collected from three retinas with a total of 16,000 stimulus presentation trials per retina. For the cells recorded and analyzed in this study, PSTHs were constructed for each stimulus color in the data set. Raster plots of the spikes across all trials and separate PSTH plots for the first and second halves of the data sets were generated and visually compared to provide an empirical index of response stationarity (Awiszus, 1997).

2.2. Estimation method

To perform statistical estimation, a forward model for neural encoding must be selected, and the model employed here begins with the non-homogenous Poisson process. This is the simplest model with the fewest assumptions for capturing a time-varying likelihood of event generation, and it uses a stochastic rate function $\lambda(t)$ as its only parameter. In this method, the rate function for each cell and stimulus was estimated by applying a unit-area Gaussian smoothing kernel (Szucs, 1998) to the PSTHs generated from training data sets. The width (σ) of this Gaussian filter and the number of training trials represent the only two free parameters for the presented method.

For a non-homogeneous Poisson process, the likelihood that a given set of observed events came from specific time-varying rate function can be directly calculated by (Snyder & Miller, 1991)

$$P(X \mid \lambda(t), t_0, t_1) = \left[\exp\left(-\int_{t_0}^{t_1} \lambda(t)\mathrm{d}t\right)\right]\left(\prod_{i=1}^{n} \lambda(x_i)\right), \quad (1)$$

where $P(X|\lambda(t), t_0, t_1)$ is the probability that the set of n events (X) over the time period t_0 to t_1 was generated by the rate function $\lambda(t)$, and $\lambda(x_i)$ is the value of the rate function at the times of the occurrence (x_i) of the n events. For time periods where no events occurred, the $\prod$ term is omitted.

For neural spike events within a trial, the stochastic rate $\lambda(t)$ over the trial is a function of the stimulus, and under the Poisson model, Eq. (1) for a single cell becomes:

$$P(X_c|s) = P(X_c \mid \lambda_c(s,t)) = \left[\exp\left(-\int_0^T \lambda_c(s,t)\mathrm{d}t\right)\right]\left(\prod_{i=1}^{n} \lambda_c(s,x_{c,i})\right), \quad (2)$$

where $P(X_c|s)$ of observing the set of spikes X_c from cell c for stimulus s, $P(X_c|\lambda_c(s,t))$ is the probability that the observed set of spikes (X_c) from cell c over the period of the trial $[0,T]$ was generated by the time-varying response function $\lambda_c(s,t)$ of cell c for stimulus s, and $\lambda_c(s,x_{c,i})$ is the value of the rate function at the times of the n spikes in the trial ($x_{c,i}$) for cell c. The intuitive interpretation of Eq. (2) is that the first term represents the penalty for not getting spikes when they are expected and the second term represents the reward for detecting spikes when likely. Although variations of the Poisson model have been widely used for modeling and estimating spike timing (see Rieke et al., 1997, for a summary), most of these focus on the likelihoods of observing specific spike counts in different bins within a trial. The continuous likelihood function shown in (2) does not require bins to be defined by the user, it only requires a description of the rate function ($\lambda(t)$) within the trial. This continuous form is recently

becoming more widely used for neural signal estimation (Brown, Barbieri, Ventura, Kass, & Frank, 2002; Kass & Ventura, 2001; Normann, Warren, Ammermuller, Fernandez, & Guillory, 2001).

The likelihood Equation (2) is optimal when neurons behave as pure Poisson processes. Real neurons, however, deviate from Poisson behavior at small time scales due to their refractory periods and these statistical deviations have been well studied (Berry & Meister, 1998; Brown et al., 2002; Johnson, 1996; Reich, Victor, & Knight, 1998). The simplest method for modeling refractory effects for a neuron is to represent its rate function $\lambda(t)$ as the product of a stimulus-driven free-firing rate $q(t)$ and a recovery function $w(t)$ which drops to 0 after a spike and returns from 0 to 1 over time. When the neural firing is controlled by the refractory periods, the PSTH deviates from the free-firing rate due to the refractory effects. Once an estimate for $w(t)$ is determined for a cell, the free firing rate function $q(t)$ can be estimated from the average PSTH rate $r(t)$ computed from M trials (Miller, 1985) by:

$$q(t) = \frac{r(t)}{\frac{1}{M}\sum_{j=1}^{M} W_j(t)}, \tag{3}$$

where $W_j(t)$ is the product of the recovery functions for all spikes over trial j. This method is also discussed in Berry and Meister (1998), along with a recovery function for the salamander retina composed of an absolute refractory period ($w(t) = 0$), followed by an exponential recovery to unity. To include the effects of the refractory period in Eq. (2), the Poisson rate function $\lambda(t)$ is defined as the estimated free firing rate $q(t)$ for the unit, multiplied by the recovery functions $w(t)$ inserted following each spike. This process is illustrated in Fig. 1A, and it effectively prevents the likelihood from being penalized for times when spikes are expected but not present due to the refractory period of the cell.

When calculating the likelihood of a pattern of spikes from a population of cells, possible dependencies or relationships in spike generation between cells must be considered. When the spikes in each cell are generated independently of the spikes of other cells, combining estimation likelihoods from the group is a simple matter of multiplication to compute the joint probability of the observed pattern for the population:

$$P(X|s) = \prod_{c=1}^{C} P(X_c|s), \tag{4}$$

where $P(X|s)$ is the probability of observing the population spike pattern X across C cells for stimulus s. Once a likelihood of observing a set of spikes given a specific stimulus s is available ($P(X|s)$), it can be used with Bayes' Rule to calculate the likelihood of a stimulus having been presented given an observed set of spikes ($P(s|X)$):

$$P(s|X) = \frac{P(X|s)P(s)}{P(X)}, \tag{5}$$

where $P(s)$ is the prior probability of stimulus s being presented, and $P(X)$ is the independent probability of the spike pattern being observed. In decoding applications in which the relative likelihoods of different stimuli are being compared for stimulus estimation, the $P(X)$ term can be considered a common normalization coefficient and ignored.

If the cells do not spike independently, it quickly becomes computationally challenging to generically model and utilize this information for estimation (Sanger, 2002). In addition, the encoding analysis and model can easily require impractical amounts of training data if too many general model terms are added to capture these dependencies and their relationships to the stimuli. However, when a clear underlying dependency model can be postulated, the dependencies can be easier to include in the analysis, especially if the underlying processes are independent. One such dependency code for the retina has been proposed by Brivanlou, Warland, and Meister (1998) in which tightly synchronized spikes between groups of closely spaced ganglion cells are driven by common amacrine cells. In this model, the amacrine cells independently code for features that are different than the features coded by the individual ganglion cells. For stimulus likelihood estimation, the amacrine and ganglion responses can be separated by removing tightly correlated (amacrine) spikes from the ganglion cell responses, and generating separate firing rate and renewal estimates for the amacrine cells. Eqs. (2)–(5) can then be applied to the ganglion and amacrine cell responses, except that the ganglion renewal functions must be applied to the ganglion rate functions for both ganglion and amacrine-induced spikes, and both the ganglion spike and amacrine spike renewal functions must be applied to the estimated amacrine rate functions.

2.3. Implementation and analysis

In this study, dependencies among pairs of recorded cells were explored using the normalized Joint PSTH (Aertsen et al., 1989; Palm, Aertsen, & Gerstein, 1988). Because of the sparse nature of the recordings, only positive correlations were investigated for significance. Based on this analysis, the cells in this study were treated as independent for stimulus estimation.

For the encoding analysis, the stochastic rate functions for each cell and stimulus were estimated by applying a unit-area Gaussian smoothing kernel to the PSTH plots. The width of the smoothing filter was optimized for stimulus estimation performance, and final analyses were performed with $\sigma = 10$ ms. For each data set that underwent decoding analysis, the stimulus estimation for each trial was performed with firing and refractory

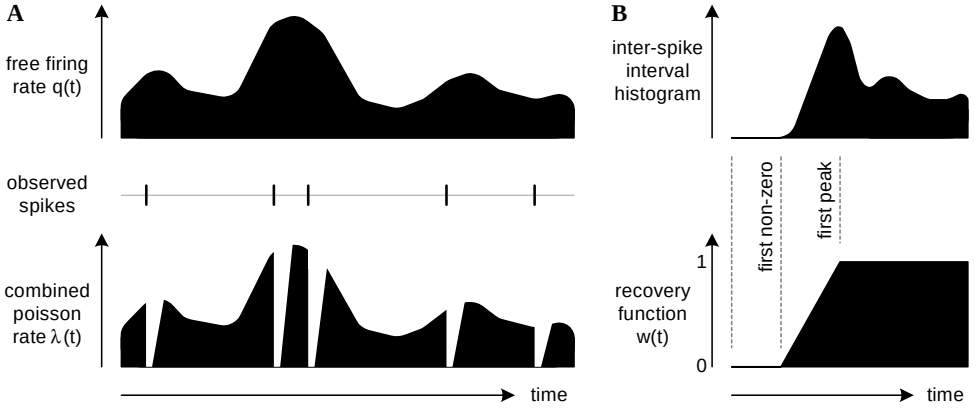

Fig. 1. (A) Computation of the combined Poisson rate function from the free firing rate and the observed spikes. The free firing rate is multiplied by the recovery function at each spike occurrence. (B) Determination of the recovery function $w(t)$ from the interspike interval histogram (ISI) of a specific channel. This function is composed of an absolute refractory period, followed by a linear recovery to unity between the times of the first non-zero and peak ISI values.

models built exclusively from all other trials in the set. For decoding analysis with models that included refractory periods, a recovery function was estimated for each cell, and the free firing rate was estimated from the PSTH according to (3). The recovery functions were difficult to estimate due to the sparse nature of the recordings and the best empirical fit for the data was provided by an absolute refractory period followed by a linear recovery to unity. The end of the absolute period was determined by the offset of the first non-zero value in the inter-spike interval (ISI) histogram, and the end of the linear recovery period was determined by the offset of the first peak in the ISI histogram. This process is illustrated in Fig. 1B. For cells that did not have strong peaks in the local ISI histogram, absolute refractory periods were used without the linear recovery phase.

For the decoding analysis, time was discretized into steps (Δt) of 1 ms for the rate functions and the likelihoods were represented by logarithms for efficient multiplication. In this implementation, Eqs. (2) and (4) become:

$$L_c(X_c|s) = \ln[P(X_c|s)] = -\left(\sum_{t=0}^{T/\Delta t} \lambda_c(s, t\cdot\Delta t)(\Delta t)\right) + \sum_{i=1}^{n} \ln[\lambda_c(s, x_{c,i})], \quad (6)$$

$$L(X|s) = \ln[p(X|s)] = \sum_{c=1}^{C} L_c(X_c|s), \quad (7)$$

where $L_c(X_c|s)$ is the log-likelihood for observing the set of spikes X_c from cell c given stimulus s, and $L(X|s)$ is the log-likelihood for the entire population. The first term of (6) serves as a first order numerical integration of the rate function, and the second term of (6) is omitted when no spikes are detected over a trial period. For each trial evaluated, stimulus likelihoods were calculated according to Bayes' Rule (5), and the stimulus with the maximum likelihood was selected as the estimate. Because only the likelihoods of stimuli relative to each other were considered, the $P(X)$ term of (5) was omitted, and the $P(s)$ term was also not included since each stimulus was equally probable within the set.

Performance for this stimulus estimation method was evaluated both with and without refractory renewal models included. The number of trials included in the analysis was optimized for stimulus estimation performance, and sets of 4000 trials were used for the final analysis. Simulated data sets were also generated for comparison to the actual retinal data. In these simulated sets, smoothed PSTH curves (Gaussian, $\sigma = 10$ ms) from actual retinal cells were used as rate functions for both Poisson generation of spikes and estimation of the spike data. These simulated spike trains were guaranteed to be stationary and free of any correlations between cells.

3. Results

Results are presented from three retinas with spikes from 18, 15, and 11 cells, respectively. Only spike recordings with signal-to-noise ratios greater than three and clear unit separation were used. In general, the ganglion cell responses were quite sparse for this preparation and mode of stimulation. For example, the most active cell had a typ-

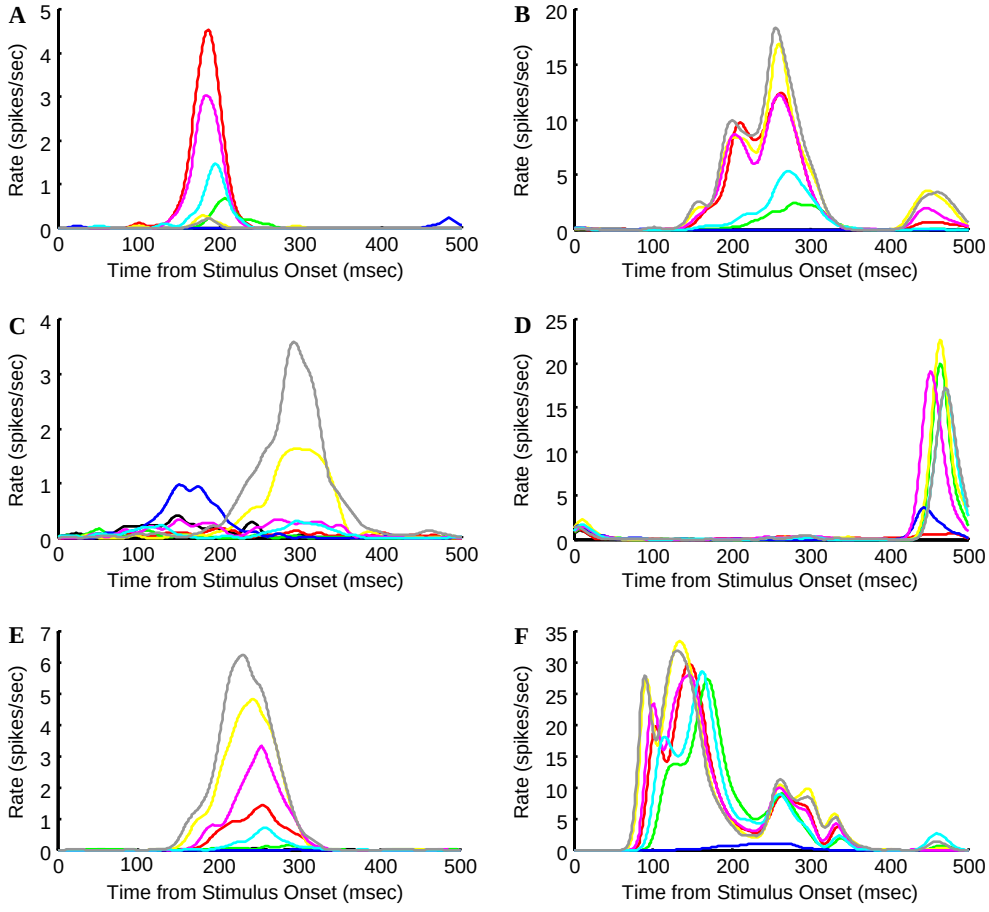

Fig. 2. Peri-Stimulus Time Histogram (PSTH) examples of general responses from recorded cells. Stimulus ON period was from 0 to 200 ms, OFF period was from 200 to 500 ms. The PSTH for each color is plotted in its stimulus color with gray substituted for white. (A) Fast ON response with strong red activation. (B) Slow ON response and slow, weak off response. (C) Slow ON–OFF response with blue/yellow opponent activation. (D) Slow-OFF response with strong red and green activation. (E) Slow ON response. (F) Fast ON–OFF response with strong red and green activation.

ical response rate of 2.3 spikes/trial, and the average for all cells was only 0.7 spikes/trial.

The responses of the cells were consistent with the classification schemes of previous research (Armington, Adolph, & Wu, 1991; Granda & Fulbrook, 1989). Representative samples of the PSTHs are shown in Fig. 2. Responses were considered fast if they contained peaks within 50–125 ms for a light change (Fig. 2F) and slow if they contained peaks with latencies greater than 125 ms (Figs. 2A, B, and E). Of the cells recorded, 66% (29/44) contained ON responses with 25% (7/29) of these showing a slow ON response. The cells with slow ON responses tended to have a more sustained response, especially to stimuli that included red. Eighty-seven percentage (39/44) of the recorded cells contained OFF responses with 56% (22/39) of these showing a slow OFF response. Thirty-six percentage (16/44) of the recorded cells showed both ON and OFF responses. The fast ON and OFF responses were typically very repeatable and often exhibited latencies that were a function of the color and intensity of the stimulus (e.g., Fig. 2F). The slower ON and OFF responses (e.g., Figs. 2A and E) had peaks that tended to overlay each other in time.

The stability of the responses over time was assessed by raster plots of the trials over the course of the experiment and Peri-Stimulus Time Histogram (PSTH) plots constructed from different periods. In the first retina, 11 of the 18 cells showed good stability, and this retina was used in the subsequent correlation and estimation analysis. The second and third retinas showed a gradual deterioration of all of the recorded cells over the course of the experiment. Figs. 3 and 4 show raster plots over the duration of the experiment for a cell with relatively good stability (3), and a cell with poor stability (4). Plots of the PSTHs over trials 4000–8000, and 10,000–14,000 from these cells are shown in Figs. 3B and C, and 4B and C. The PSTHs of the cell in Fig. 4 showed a drift from slow ON and OFF responses to a faster OFF response, and similar changes were seen in many of the nonstationary recorded cells. It was unclear whether this was a natural change of some cells in the retina due to the stimulation or changes induced by the invasive nature of this preparation and extracellular recording technique. The effects of these drifts on the performance of the estimation as more trials are used for the model building are discussed below.

The statistical dependencies between the cells were assessed with Joint peristimulus time histogram (JPSTH) analysis. Fig. 5 shows the JPSTH plots of the cell pair with the highest correlations found. Fig. 5A shows a 2D histogram of the coincident firing activity (cross-product of the spike trains of a trial for two units) across all stimulus trials. This plot represents the observed coincident firing rates $P(a*b)$ in the data. Fig. 5B represents the predicted coincident firing rate $P(a)*P(b)$ given an assumption of independence and Fig. 5C represents the difference between Figs. 5A and B. Fig. 5D shows the 1 ms bins in Fig. 5C that are above a 97.5% likelihood of being due to chance. In this plot, significant correlations show up as dense areas, and collapsing the plot along the diagonal produces a simple cross-correlogram with the stimulus-related correlation removed. For all cell pairs, including the pair shown in Fig. 5D, the bins were sparse, unstructured, and not indicative of correlations between the cells.

The stimulus estimation presented here has two free variables: the number of trials included in the model building and analysis set, and the width of the Gaussian smoothing filters used to build the response probabilities from the PSTH data. Estimation performance for all 18 cells in the

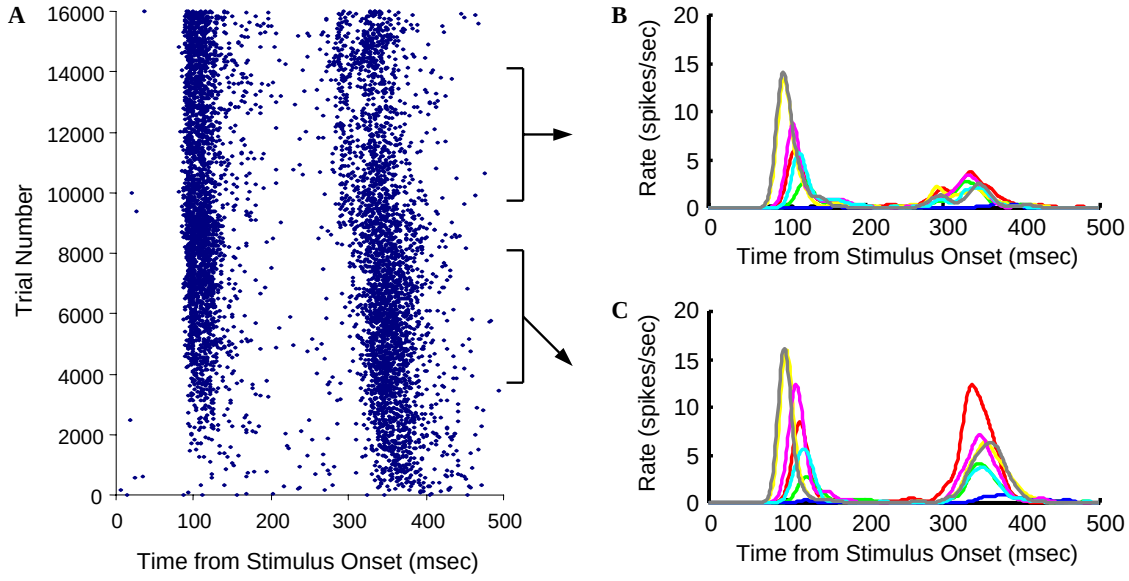

Fig. 3. (A) Raster Plot from 16,000 trials of a cell with relatively good stationarity. (B) PSTH of the cell for trials 10,000–14,000 for each stimulus color. (C) PSTH of the cell for trials 4000–8000 for each stimulus color. Stimulus ON period was 0–200 ms, OFF period was 200–500 ms. Although there are significant changes in the OFF response component throughout the experiment, the ON response is relatively consistent.

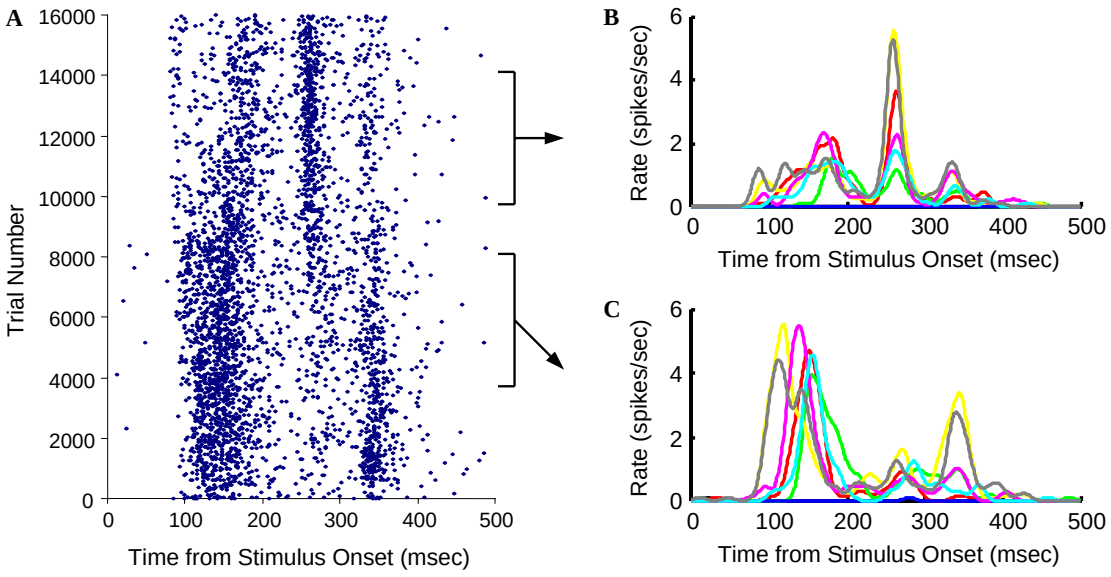

Fig. 4. (A) Raster Plot from 16,000 trials of a cell with relatively poor stationarity. (B) PSTH of the cell for trials 10,000–14,000 for each stimulus color. (C) PSTH of the cell for trials 4000–8000 for each stimulus color. Stimulus ON period was 0–200 ms, OFF period was 200–500 ms. Note the changes in response characteristics between the first and second half of the experiment.

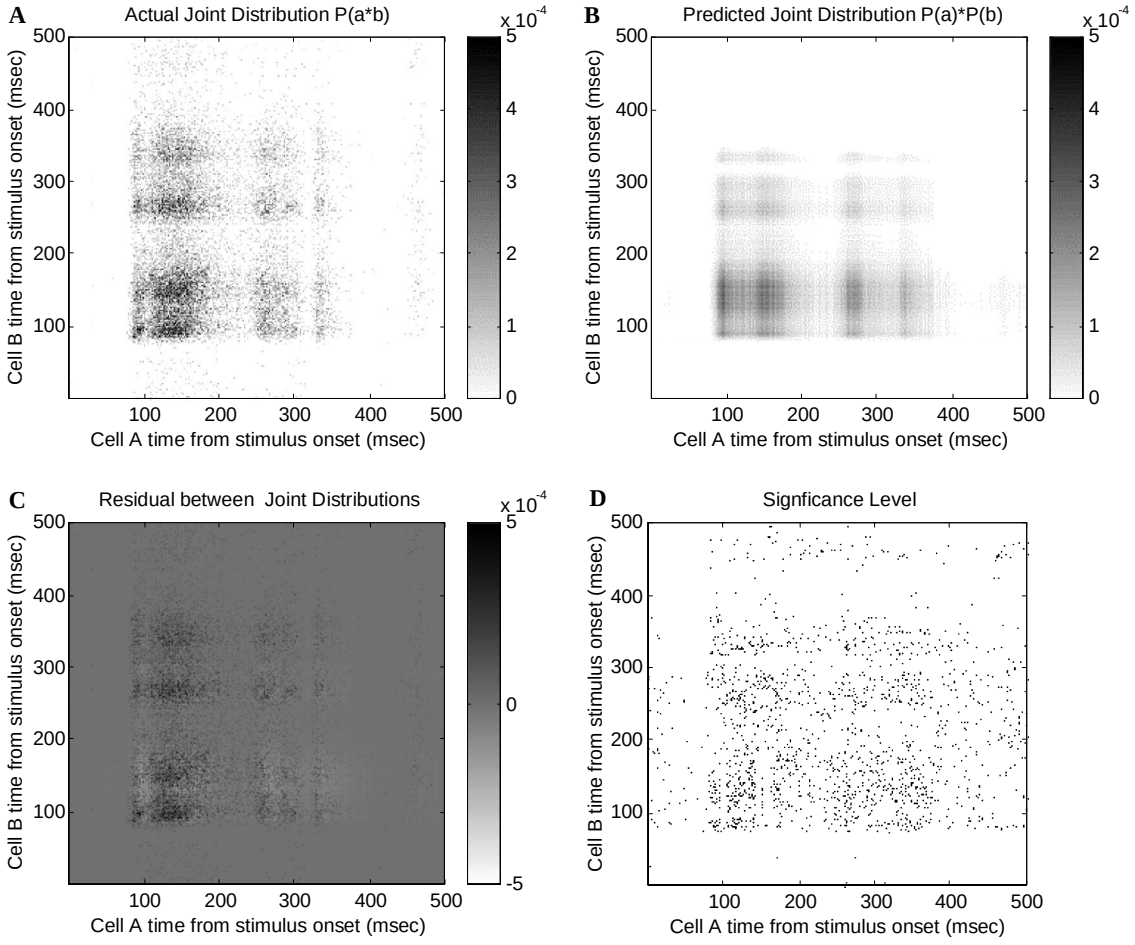

Fig. 5. Joint Peri-Stimulus Time Histograms (JPSTH) analysis of neighboring cells with the strongest correlations in the data set. (A) 2D histogram of actual coincident firing probabilities. (B) Plot of predicted coincident firing rates from the individual PSTH responses. (C) Differences between the recorded and predicted coincident responses. (D) Plot of bins with coincident firing that have a probability above chance of >0.975.

first retina is plotted as a function of these two variables in Fig. 6. Smoothing the PSTH generalizes the response model obtained from the available data, but too much can overly blur the distinctions between the responses. Likewise, more trials are better for model building until the non-stationary drift in the response characteristics begins to limit the accuracy and performance of the models. The peak performances were in the range of 7–12 ms for the filter widths (σ) and approximately 3000–4500 trials (a time window of 25–38 min). Especially poor performance resulted from using less than 1000 trials and smoothing filters narrower than 3 ms.

Final stimulus decoding estimations were generated for the 18 cells using 4000 trials and PSTH smoothing filters with $\sigma = 10$ ms. The performances as a function of the number of cells are shown in Fig. 7A. This plot shows the accuracy for the correct stimulus being estimated by the maximum likelihood as well as the performance for the correct stimulus being in the top two and top three likeliest values. The cells were ordered to provide the quickest convergence to the maximum value. Table 1 shows the estimation performance for all 18 cells broken down by individual color. The highest average performance for this estimation task was 78% as compared with 12.5% for chance. Estimation for the models that included individual cell refractory periods performed nearly identically (within 1% at each cell count) to the models and estimation functions without cell refractory periods.

Fig. 7B shows the performance for the 18 recorded cells versus simulated spike trains generated from the smoothed PSTHs of these cells according to an ideal Poisson model. The performance curves for the actual and simulated cells were nearly identical for up to 11 cells, and the simulated cells performed incrementally better with 12–18 cells. This similarity in performance suggests that, given the encoding model, the estimation performance was maximized for the available set of population recordings. This difference at the end of the curve is likely due to the perfect stationarity of the simulated cells compared to the actual cells. Combinations of multiple, independent simulated responses from

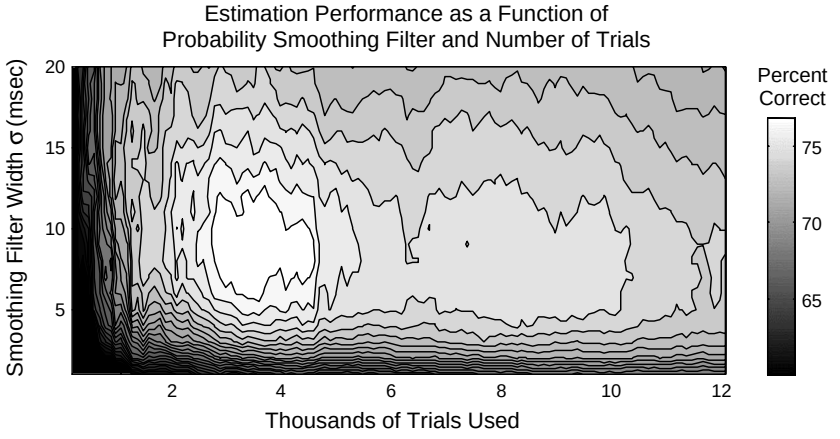

Fig. 6. Estimation performance as a function of the number of trials included in the estimation and the standard deviation of the smoothing filter used to estimate firing probabilities from the peri-stimulus time histograms (PSTH). Gray intensities are scaled to indicated performances between 60%and 78% correct.

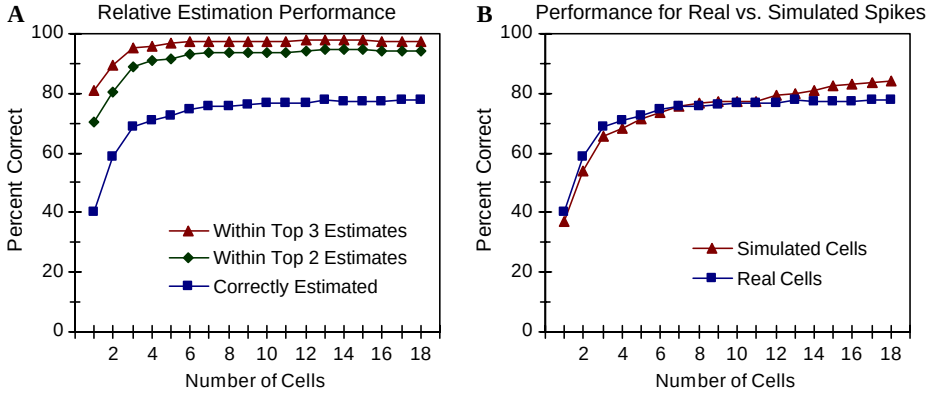

Fig. 7. (A) Estimation performance as a function of the number of cells used. The order of cells was selected to provide the fastest rise in accuracy. Plots are shown for the correct value being the most likely according to the estimation, and within the top two and top three likeliest values. (B) Performance for correct estimation for real cells and simulated cells.

Table 1
Estimation performance for different stimulus colors

	Blk	Red	Grn	Yel	Blu	Cyn	Mag	Whi	Total
Total presented	491	468	509	541	494	494	512	491	4000
Total guessed	489	449	582	546	473	484	488	489	4000
Estimated correctly (%)	95	87	74	69	89	77	64	70	78
Within top two (%)	97	95	94	95	95	89	93	97	94
Within top three (%)	97	97	98	98	96	98	97	99	98

The total number of trials presented for each color is shown in the top of the table along with the number of times the color was estimated as the most likely stimulus presented. Individual stimulus performances are also shown for the correct color being within the top two and top three most likely stimulus.

the best cells in the set (such as the cell in Fig. 2F) performed better with up to 90% correct classification for 10 cells.

4. Discussion

Once a stochastic model for neural encoding has been proposed, an optimal decoding method can almost always be directly derived. The method presented here for estimating flashed visual stimuli based on the neural response starts with one of the simplest encoding models, a non-homogeneous Poisson rate code that varies as a function of stimulus. This model captures the intuition associated with the PSTH—that there is an average response rate that varies in time for each stimulus. The primary advantages of this method over many other trial-based decoding techniques are the lack of subjective parameters such as bin sizes and locations and the ability to incorporate refractory behavior. The two free variables for the method, the width of the model smoothing filter and the number of trials used to build the response models, can be easily investigated to find their optimum values based on the application or data. This allows the Poisson model with refractory period to be used as a standard or starting point for evaluating other neural coding models and decoding methods for trial-based experiments.

All of the recorded cells showed some nonstationary response changes over the course of the experiment. For stochastic estimation, these changes limit the range of time that can be used for model building and this experiment reached a maximum performance using approximately 4000 trials or 30 min of data. Alternatively, one can use models that vary in time or adaptively evolve based on recent responses (Brown, Nguyen, Frank, Wilson, & Solo, 2001; Johnson, 1996). For estimation studies, these changes also highlight the importance of using randomly shuffled stimulus sets rather than each individual stimulus in series.

The responses for each trial typically consisted of only a few spikes across all electrodes, and the firing did not appear to be limited by refractory periods. This sparseness and the high percentage of cells with OFF responses was typical of the recordings from this preparation and recording array. It was therefore not surprising to find an only marginal (typically less than 1%) improvement in performance, when the refractory recovery functions were included in the model building and stimulus estimation. Experiments with higher firing rates are likely to be required for the refractory inclusion to have a significant contribution.

The presence of correlations and possible synchrony codes in the salamander retina has been reported for electrodes closer than 400 μm (Brivanlou et al., 1998; Meister et al., 1995). The prevalence and significance of such synchrony codes in vertebrates is still a subject of research and there is some evidence that ganglion cells act as relatively independent information encoders in some species (Chichilnisky & Baylor, 1999; Nirenberg, Carcieri, Jacobs, & Latham, 2001; Latham & Nirenberg, 2005). The electrode spacing in the array used in this study was 400 μm and no such correlations were found.

Over the last couple of decades there has been a constant dialog over the relative importance of precise spike timing codes versus rate codes in the nervous system (Softky, 1995). However, both of these models can be captured by a stochastic process with a rate function that varies in time (Berry & Meister, 1998). Precise timing codes are reflected by sharp peaks in the rate function, whereas rate codes are characterized by smoother, constant rate values. Most of the cells analyzed in these experiments exhibited both rate and precise timing characteristics in response to the flashed stimulus. For example, the cell in Fig. 2F shows different latencies for the first and second spikes in the response for each stimulus. This cell also shows differences in the areas under the response curves (or average rates) for each stimulus. This combination of specific spike timing and local spike rate differences for each stimulus can also be seen in the cells in Figs. 2A, D, and E. The stochastic estimation method presented here provides a convenient way to capture both of these aspects without having to set up special bin structures and other metrics such as the response latencies of the first and second spikes.

When stimulus sets are arbitrarily chosen, the estimation performance from a decoding algorithm becomes somewhat meaningless since any type of easy or difficult set can be used. The eight stimuli used in this study formed a very simple closed set that was not intended to exhaustively explore the space of visual stimuli. Instead, this set was chosen to provide an illustrative example of how the proposed estimation method can be used for trial-based estimation tasks. In general, neural decoding research

should always focus on model-based approaches that include the definition of the encoding model and the derivation of the optimal inverse estimation method for the model. An understanding of the forward and inverse models for any neural system allows the fundamental limits for coding and decoding performance to be determined, and it allows the neural system under study to be placed in the context of larger neural systems.

When more information is known about the coding structure of a neural system, more sophisticated decoding models can be constructed to utilize this knowledge. Recent likelihood decoding methods have been created for hippocampal place cells (Brown et al., 2001) and motor cortex (Barbieri et al., 2004; Brockwell, Rojas, & Kass, 2004; Kemere, Shenoy, & Meng, 2004; Shoham, 2001; Shoham et al., 2005; Wu et al., 2004). In addition, improved methods have recently been published for estimation of neural firing rates from the PSTH based on likelihood (Kass, Ventura, & Cai, 2003), and for evaluating the fit of a stochastic forward model to the spike data (Brown et al., 2002). Future decoding studies for the visual system should consider these recent works, as well as focus on ways to use these techniques for estimation of continuous stimuli and natural visual scenes.

Acknowledgments

This work funded by NSF Grant IBN-9424509 and a Utah State Center of Excellence grant. All animal procedures were performed under the review of the University of Utah Institutional Animal Care and Use Committee (IACUC) and in accord with federal regulations. Special thanks to Vinh Ngo and Yan-Ping Li at the Utah Center for Neural Interfaces for their assistance with electrode fabrication and assembly.

References

Abeles, M. (1991). *Corticonics: neural circuits of the cerebral cortex*. Cambridge: Cambridge University Press.

Aertsen, A. M. H. J., Gerstein, G. L., Habib, M. K., & Palm, G. (1989). Dynamics of neuronal firing correlation: modulation of effective connectivity. *Journal of Neurophysiology, 61*, 900–917.

Armington, J. C., Adolph, A. R., & Wu, S. (1991). Spatial properties of ganglion cell activity in the turtle retina. *Visual Neuroscience, 6*, 439–450.

Awiszus, F. (1997). Spike train analysis. *Journal of Neuroscience Methods, 74*, 155–166.

Barbieri, R., Frank, L. M., Nguyen, D. P., Quirk, M. C., Solo, V., Wilson, M. A., et al. (2004). Dynamic analyses of information encoding in neural ensembles. *Neural Computation, 16*(2), 277–307.

Becker, J. D., & Kruger, J. (1996). Recognition of visual stimuli from multiple neuronal activity in monkey visual cortex. *Biological Cybernetics, 74*, 287–298.

Berry, M. J., & Meister, M. (1998). Refractoriness and neural precision. *Journal of Neuroscience, 18*(6), 2200–2211.

Brivanlou, I. H., Warland, D. K., & Meister, M. (1998). Mechanisms of concerted firing among retinal ganglion cells. *Neuron, 20*, 527–539.

Brockwell, A. E., Rojas, A. L., & Kass, R. E. (2004). Recursive bayesian decoding of motor cortical signals by particle filtering. *Journal of Neurophysiology, 91*(4), 1899–1907.

Brown, E. N., Nguyen, D. P., Frank, L. M., Wilson, M. A., & Solo, V. (2001). An analysis of neural receptive field plasticity by point process adaptive filtering. *Proceedings of the National Academy of Science USA, 98*(21), 12261–12266.

Brown, E. N., Barbieri, R., Ventura, V., Kass, R. E., & Frank, L. M. (2002). The time-rescaling theorem and its application to neural spike train data analysis. *Neural Computation, 14*(2), 325–346.

Brown, E. N., Barbieri, R., Eden, U. T., & Frank, L. M. (2003). Likelihood methods for neural data analysis. In J. Feng (Ed.), *Computational neuroscience: A comprehensive approach* (pp. 253–286). London: CRC.

Chichilnisky, E. J., & Baylor, D. A. (1999). Synchronized firing by ganglion cells in monkey retina. *Annual Meeting of the Society for Neuroscience*, Abstract 424.9.

Gawne, T. J., & Richmond, B. J. (1993). How independent are the messages carried by adjacent inferior temporal cortical neurons?. *Journal of Neuroscience 13*(7), 2758–2771.

Geisler, W. S., & Albrecht, D. G. (1997). Visual cortex neurons in monkeys and cats: detection, discrimination, and identification. *Visual Neuroscience, 14*, 897–919.

Granda, A. M., & Fulbrook, J. E. (1989). Classification of turtle retinal ganglion cells. *Journal of Neurophysiology, 62*, 723–737.

Guillory, K. S., & Normann, R. A. (1999). A 100-channel system for real time detection and storage of extracellular spike waveforms. *Journal of Neuroscience Methods, 91*(1-2), 21–29.

Hartline, H. K. (1938). The response of single optic nerve fibers of the vertebrate eye to illumination of the retina. *American Journal of Physiology, 121*, 400–415.

Johnson, D. H. (1996). Point process models of single-neuron discharges. *Journal of Computational Neuroscience, 3*, 275–299.

Kass, R. E., & Ventura, V. (2001). A spike-train probability model. *Neural Computation, 13*(8), 1713–1720.

Kass, R. E., Ventura, V., & Cai, C. (2003). Statistical smoothing of neuronal data. *Network, 14*(1), 5–15.

Kemere, C., Shenoy, K. V., & Meng, T. H. (2004). Model-based neural decoding of reaching movements: A maximum likelihood approach. *IEEE Transations on Bio-medical Engineering, 51*(6), 925–932.

Kuffler, S. W. (1953). Discharge patterns and functional organization of mammalian retina. *Journal of Neurophysiology, 16*, 37–68.

Latham, P. E., & Nirenberg, S. (2005). Synergy, redundancy, and independence in population codes, revisited. *Journal of Neuroscience, 25*(21), 5195–5206.

Meister, M., Lagnado, L., & Baylor, D. A. (1995). Concerted signaling by retinal ganglion cells. *Science, 270*, 1207–1210.

Miller, M. I. (1985). Algorithms for removing recovery-related distortion from auditory-nerve discharge patterns. *Journal of the Acoutical Society of America, 77*, 1452–1464.

Nirenberg, S., Carcieri, S. M., Jacobs, A. L., & Latham, P. E. (2001). Retinal ganglion cells act largely as independent encoders. *Nature, 411*(6838), 698–701.

Normann, R. A., Warren, D. J., Ammermuller, J., Fernandez, E., & Guillory, S. (2001). High-resolution spatio-temporal mapping of visual pathways using multi-electrode arrays. *Vision Research, 41*(10-11), 1261–1275.

Oram, M. W., Foldiak, P., Perrett, D. I., & Sengpiel, F. (1998). The 'Ideal Homunculus': Decoding neural population signals. *Trends in Neuroscience, 21*, 259–265.

Palm, G., Aertsen, A. M. H. J., & Gerstein, G. L. (1988). On the significance of correlations among neuronal spike trains. *Biological Cybernetics, 59*, 1–11.

Perkel, D. H., Gerstein, G. L., & Moore, G. P. (1967). Neuronal spike trains and stochastic point processes: II simultaneous spike trains. *Biophysical Journal, 7*(4), 419–440.

Perlman, I., Normann, R. A., Chandler, J. P., & Lipetz, L. E. (1990). Effects of calcium ions on L-type horizontal cells in the isolated turtle retina. *Visual Neuroscience, 4*(1), 53–62.

Reich, D. S., Victor, J. D., & Knight, B. W. (1998). The power ratio and the interval map: Spiking models and extracellular recordings. *Journal of Neuroscience, 18*(23), 10090–10104.

Reich, D. S., Victor, J. D., Knight, B. W., Ozaki, T., & Kaplan, E. (1997). Response variability and timing precision of neuronal spike trains in vivo. *Journal of Neurophysiology, 77*(5), 2836–2841.

Rieke, F., Warland, D., de Ruyter van Steveninck, R., & Bialek, W. (1997). *Spikes: Exploring the neural code*. Cambridge, MA: MIT Press.

Salinas, E., & Abbott, L. F. (1994). Vector reconstruction from firing rates. *Journal of Computational Neuroscience, 1*, 89–107.

Sanger, T. D. (2002). Decoding neural spike trains: Calculating the probability that a spike train and an external signal are related. *Journal of Neurophysiology, 87*(3), 1659–1663.

Shoham, S., Paninski, L. M., Fellows, M. R., Hatsopoulos, N. G., Donoghue, J. P., & Normann, R. A. (2005). Statistical encoding model for a primary motor cortical brain-machine interface. *IEEE Transactions in Biomedical Engineering, 52*(7), 1312–1322.

Shoham, S. (2001). Advances towards an implantable motor cortical interface. Ph.D. dissertation. Department of Bioengineering, University of Utah.

Shoham, S., Fellows, M. R., & Normann, R. A. (2003). Robust, automatic spike sorting using mixtures of multivariate *t*-distributions. *Journal of Neuroscience Methods, 127*(2), 111–122.

Softky, W. R. (1995). Simple codes versus efficient codes. *Current Opinion in Neurobiology, 5*, 239–247.

Snyder, D. L., & Miller, M. I. (1991). *Random point processes in time and space* (second ed.). New York: Springer-Verlag.

Szucs, A. (1998). Applications of the spike density function in analysis of neuronal firing patterns. *Journal of Neuroscience Methods, 81*, 159–167.

Wu, W., Black, M. J., Mumford, D., Gao, Y., Bienenstock, E., & Donoghue, J. P. (2004). Modeling and decoding motor cortical activity using a switching Kalman filter. *IEEE Transactions of Bio-medical Engineering*(6), 933–942.