Markov Models for Neuronal Spike Trains

Ph.D. Thesis Proposal

Jeffrey Liebner
Department of Statistics, Carnegie Mellon University

ABSTRACT

Sequences of action potentials from neurons, known as spike trains, can be modeled as point processes. Generally, the firing behavior of a neuron is dependent upon the past spiking history of the neuron, a behavior that can be described using Markov models. Also, neurons may exhibit variations during repeated trials of an experiment. In addition, the firing patterns of neurons may be dependent upon outside sources, such as oscillatory behavior in the brain or the firing of neighboring neurons. The goal of this dissertation is to propose a model that incorporates each of these behaviors in a Bayesian framework. A Gibbs sampling procedure is described which addresses each neuronal effect in turn and utilizes the mechanism of Bayesian Adaptive Regression Splines (BARS). The performance of the procedure will be evaluated through simulation, and then the apparatus will be applied to actual neuronal data. Finally, the BARS mechanism will be addressed, evaluating its coverage adequacy.

1 Introduction

Neurons communicate through sequences of action potentials, known as spike trains. In probability theory and statistics, a time series of discrete events, such as a spike train, is called a point process (Daley and Vere-Jones, 2003). In general, point processes allow the probability of a spike at time t to depend in arbitrary ways on the history of the spike time prior to time t. Markov models provide a useful simplification. This thesis will propose markov models that address the temporal dependence of an individual neuron's firing behavior on its past spiking history while also modeling variations among repetitions (trials) of experiments.

Figure 1 displays a set of data recorded from a single neuron during 15 trials of a single set of experimental conditions. In this case, as in other experimental studies, there is an event that defines the time $t = 0$. Here, the data are recorded from the olfactory bulb of a locust, and the event that was labeled $t = 0$ is the exposure of the insect to the scent hexanol, with exposure lasting for one second. This figure shows the response of the neuron two to four seconds after the initial exposure of the scent. The left figure is a raster plot. In this plot, each spiking occurrence is displayed as a vertical slash, and the results of successive trials are stacked on top of each other. The right figure is the peristimulus time histogram (PSTH). This is obtained by pooling the spiking times from the separate trials into a single histogram.

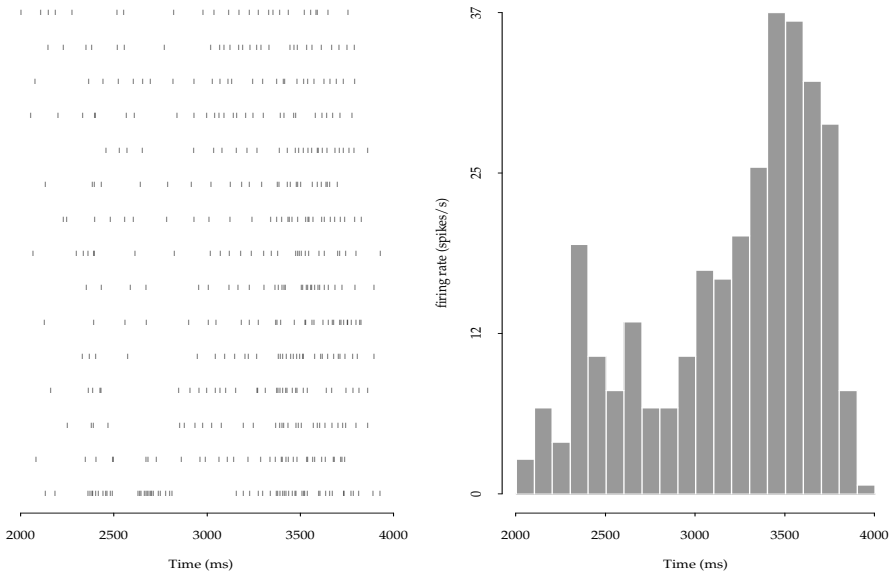

Figure 1: Raster plot and peristimulus time histogram (PSTH) for a single neuron recorded in the olfactory bulb of a locust. In the figure, $t = 0$ marks the time the odor stimulus, here hexanol, was released. LEFT: The raster plot of observed firing times for each of 15 experimental trials, with the first trial on the bottom of the plot. Each tick mark represents the timing of a spike recorded from the neuron. RIGHT: The PSTH, which displays firing counts occurring in 100 millisecond bins, normalized to units of spikes per second.

Since the neurons communicate with each other through their spiking mechanism, examination of these spike trains can reveal information about the processes occurring in the brain. Studying the relationships between neural firing patterns and behavior can help lead to an understanding of how the brain uses and stores information. In order to conceptualize the mechanisms that underlie information coding and decoding, developing models that describe the spiking behavior is necessary.

The data displayed in figure 1 can be represented by letting $s_1, s_2, ...S_n$ represent the spike times up to time t and H_t denote the spiking history up to time t. In an infinitesimal unit of time, the probability of

spiking for a Poisson process is denoted as $\lambda(t)dt$, where $\lambda(t)$ is referred to as the time-varying rate, or the intensity function. For a general point process, the intensity function is conditional upon the history, and is denoted $\lambda(t|H_t)$. Thus, the probability of a spike in $(t, t+dt)$ is $\lambda(t|H_t)dt$.

In practice, for small Δt the point process likelihood is approximated by the corresponding binary likelihood, written in terms of $p(t_i|H_{t_i}) = P(X(t_i) = 1|H_{t_i})$, the probability of observing a spike at time t_i. Here the data were represented as a binary sequence $X(t_1), X(t_2), ..., X(t_{max})$ with $X(t_j) = 1$ signifying that a spike has occurred in the time bin $(t_j - \Delta t/2, t_j + \Delta t/2)$ and $X(t_j) = 0$ otherwise. Using the binary representation of the data, an approximation to the probability of the spike train during the time interval $[0, T]$ is

$$p(s_1, ...s_n) = \frac{1}{(\Delta t)^n} * \prod_{i=1} p(t_i|H_{t_i})^{x(t_i)}(1 - p(t_i|H_{t_i}))^{1-x(t_i)}. \quad (1)$$

If one takes the limit of this expression as $\Delta t \to 0$, the following expression is obtained

$$p(s_1, ...s_n) = e^{-\int_0^T \lambda(t|H_t)dt} \prod_{k=1} \lambda(s_k|H_{s_k}). \quad (2)$$

As Brillinger (1988, 1992) observed, this result allows generalized linear and nonlinear modeling strategies to be effective in analyzing point process data.

One of the simplest ways to model spike times is to assume that they follow a Poisson process. In the Poisson case, the conditional intensity is equal to the unconditional or marginal intensity function $\lambda(t)$. Due to a general limit theorem (Daley and Jones, 1988, Theorem 9.2V), combined identical independently distributed point processes will converge to a Poisson process as the number of trials becomes infinite. Thus, the Poisson model often provides an adequate description of spike data pooled over many trials. However, this assumption is usually insufficient when modeling within-trial firing behavior. In developing the main model for this thesis, several deviations from independent identically distributed Poisson process trials will be incorporated.

1. During repeated trials of an experiment, the data may exhibit trial-dependent, slowly-varying changes in firing rate, which produce excess trial-to-trial variation, beyond what would be expected by random chance. Two of the trial dependent effects are latency and excitability (Brody, 1999a, 1999b; Baker and Gerstein, 2001). Latency refers to a variation in the time of onset of response of a neuron to a stimulus, a shift in time of the response. This would lead one to realign the trials based on an estimate of the time shift (Woody, 1967; Baker and Gerstein, 2001; Ventura, 2004). Excitability describes a change in the amplitude of the firing rate over repeated experimental trials. Efforts to model this change in firing rate have also been performed (Brody, 1999a, 1999b; Nawrot, et al., 1999, Ventura, et al., 2005b). These effects will be discussed later.

2. The inherent dynamics of the cell have been shown to produce non-Poisson patterns even in the absence of rate modulation (Gabbiani and Koch, 1998; Reich, et al., 1998). Neurons have a refractory period during which, after each spike, the probability of firing is greatly reduced. In addition, some neurons exhibit bursting behavior, short periods of time during which the firing rate of the neuron is extremely high. Furthermore, the simplest of dynamical models of neurons result in probability models for the interspike interval distribution that are right-skewed and non-exponential, often approximated better by an inverse Gaussian distribution (Gerstein and Mandelbrot, 1964).

3. Measurements of activity have shown underlying oscillatory behavior, for example the theta rhythm observed in the hippocampus (e.g. Hasselmo, et al., 2002). Research has suggested that certain

neurons become tuned to this oscillatory behavior, and the firing rate of the individual neurons is directly related to an aspect of these oscillations (Stopfer and Laurent, 1999).

4. The firing of other neurons either directly or indirectly connected to that neuron may cause an increase or decrease in the firing rate. Certain neurons may be more likely to fire within a certain time interval of each other, or may inhibit their firing. Certain models such as those developed by Truccolo, et al. (2005) attempt to incorporate these ensemble effects into the conditional intensity function for an individual neuron.

The model proposed for this thesis incorporates each of these deviations from the basic Poisson process model. Bayesian estimation via MCMC produces uncertainty measures associated with each of these deviations.

2 Background

2.1 Motivation

The task of neural science is to explain behavior in terms of the activities of the brain. Of particular interest is how individual nerve cells perform and interact in order to produce behavior and how these individual cells are influenced by their environment. The challenge is to understand the mental processes by which an individual perceives, acts, learns, and remembers (Kandel, et al., 2000). All animals need to know what is happening in the surrounding environment. Thus, the brain has developed mechanisms to gather and organize sensory information, creating both temporary and permanent representations of this information. The research of Mark Stopfer at the National Institutes of Health and Gilles Laurent at the California Institute of Technology has focused on how neural systems process sensory information, specifically that of olfactory input. Using the olfactory system as a model, they have examined neural coding and circuit dynamics in the brain, seeking to develop insights into the mechanisms by which information is received and stored.

The work of Stopfer and Laurent has incorporated electrophysiological, imaging, genetic, and behavioral techniques, typically applied to animals with relatively small brains, such as locusts and honeybees, with the hope of discovering functional principles of widespread relevance. For example, the data displayed in Figure 1 were obtained by exposing a locust to an odor, in this particular case hexanol, and then recording the brain activity of projection neurons in the olfactory bulb of the locust. Different concentrations of the odor, as well as chemically similar but different odors, were used to examine the change in neural response to different sensory inputs. This was performed to examine the mechanisms that underlie the phenomenon of stimulus invariance, in which stimuli are perceived as the same, or of the same class, despite actual differences in orientation or intensity. These mechanisms have been studied through repeated experimentation, allowing for examination of the changes in the response of the neuron over time.

In their experiments, Stopfer and Laurent (e.g. 1999) have observed several changes in the effects within their experiments with locusts, such as observing a significant decrease in the firing activity of neurons in the olfactory bulb following repeated exposure to stimuli. Neuronal activity dropped from the first trial of an experiment to the second. After the second trial, the firing intensity of individual neurons remained fairly constant over repeated trials. In addition, scientists have observed that certain measurements of brain activity demonstrate oscillatory effects (e.g. Hasselmo, et al., 2002; Laurent, 2002). Examples of oscillatory behavior have been shown to be related to the firing behavior of individual neurons. A local field potential (LFP) is the potential caused by the sum of all dendritic synaptic activity within a volume of tissue and has the property of often exhibiting oscillatory behavior. In the olfaction studies, the coherence of the LFP waveform recorded in the mushroom body of locusts has been shown to increase by a factor of 4 to 5 over the course of an experiment (Stopfer and Laurent, 1999). This development of a oscillation in the LFP at a

particular frequency has been used to argue for mechanisms of learning. Current research (Stopfer, et al., 2003) suggests that the instantaneous phase of the oscillatory process, such as LFP, is related to the firing rate of individual neurons. Regardless of odor concentration, individual neurons were more likely to fire at a particular phase of the LFP signal. Their studies have allowed them to establish the functional relevance of emergent circuit properties in the oscillatory system.

While the work of Laurent and Stopfer has revealed promising results, certain aspects of their published statistical results present opportunities for improvement. Changes in firing activity over the course of successive trials was quanitified using overall spike counts during a whole trial, as opposed to a model which incorporated experimental time as well. While efforts were taken to analyze the development of a coherent oscillatory signal in the LFP, only rudimentary techniques were used to associate the LFP to neuronal firing. Auto-correlograms were used to find oscillatory nature in the firing of individual neurons, and voltage tracings of LFP and individual neurons were compared by eye to determine if the neurons appeared to fire at particular phases of the LFP. Similar comparisons were used to describe the synchronization of multiple neurons to each other (Stopfer and Laurent, 1999). Later efforts used a registration point method to estimate phase by estimating the time of a spike relative to a peak or trough in the LFP signal (Stopfer, et al., 2003). While this analysis showed increases in spiking activity at particular phases of the LFP, the methodology is rather crude and doesn't include measures of uncertainty in the estimation of the phase. In addition, the focus of this work was on the effect of varying concentrations of odor stimuli upon neuron activity rather than adaptation across time.

In order to facilitate the research, this work will propose a model and methodology that is able to quantify and assess the significance of changes in the activity of individual neurons. The model will be able to incorporate information about variation across trials and information about outside oscillatory behavior while also including experimental time.

2.2 Trial-to-trial variability

Spike trains recorded from animals display variation in spike timing both within and across repeated trials. In some cases, the variation may be considered noise. However, trial-to-trial spike train variation that is greater than what would be expected by random variation may be of interest for its physiological significance and for its effect on statistical procedures.

2.2.1 Latency

One form of trial-to-trial variability is known as latency. With latency, there is a tendency for the neuronal response to shift in time, or that it might fire earlier or later on some trials than on others. Thus, the firing rate of trial k may be represented by $\lambda(t - \tau_k)$ where τ_k represents the time lag for trial k.

A few methods for estimating the latency have been developed. Brody (1999a) proposed the following brute force method. Several possible equally spaced estimates of the latency τ_k were evaluated for each trial, searching for the minimum of a pre-defined cost function. After the cost-function was minimized for each trial, the process was looped several times until a stable estimate was obtained.

Another method developed by Ventura (2004) is much less intensive computationally. It relies on the result that the spike times of a Poisson process with firing rate of $\lambda(t)$ may be viewed as a random sample from a distribution with density proportional to $\lambda(t)$. The spike times may be combined over all K trials and viewed as a random sample with a distribution that is proportional to

$$\sum_{k=1}^{K} \lambda(t - \tau_k).$$

This distribution has the same basic shape as the PSTH before adjusting for latencies. The latencies τ_k are chosen so that the variance of the PSTH is minimized.

2.2.2 Excitability

Excitability describes variation in the amplitude and/or shape of the firing rate. It is possible to model the fluctuation of the firing rate of individual trials using kernel smoothers or by using the interspike intervals (ISIs). However, the spike train for an individual trial may be fairly sparse, and these techniques ignore information regarding the overall firing rate of the neuron obtained from the aggregated trials, summarized by the smoothed PSTH.

To address the issue of trial-to-trial excitability, Ventura, et al. (2005b) proposed an alternative method. First, a smoothed estimate $\hat{\lambda}(t)$ of the overall firing rate $\lambda(t)$ was obtained. The firing rate of the individual experimental trials was also estimated, using regression splines. Next, an estimate of the deviation of the firing rate of each trial k relative to the overall firing rate of the neuron was determined with the following ratio:

$$\bar{g}_k(t) = \hat{\lambda}^k(t)/\hat{\lambda}(t)$$

with $\lambda^k(t)$ representing the estimated firing rate of trial k. Other methods, such as that of Brody (1999a), insisted that the function $\bar{g}_k$ be a constant.

Because the data for individual trials tends to be sparse, the deviations for $\bar{g}_k(t)$ tend to be highly variable. One solution would be to smooth the deviations $\bar{g}_k(t)$ with splines. However, to prevent simply fitting noise, the splines should be fit using very few ideally placed knots. This would allow the trial-to-trial excitability fit to be curvilinear but constrained to be slowly varying across time. Ventura, et al.. (2005b) used two equally spaced knots to perform the smoothing, but ideally one would want to use knots that were adaptive to the data, instead of being fixed.

2.3 Markov Models

When examining the spike train of a neuron, it often becomes immediately clear that the firing behavior of a neuron is related to its past firing history. As mentioned earlier, this can be related to the inherent dynamics of the cell itself, including factors such as refractory period. This observation leads one to believe that an appropriate model addressing this deviation from Poisson behavior should incorporate a history component.

One model that incorporates this idea of a history component is an autoregressive model, such as that proposed by Smith and Brown (2003). In this model, a latent process modulating the neural spiking activity was represented as a gaussian autoregressive process. Given the latent process, the neural spiking activity is characterized as a general point process defined by its conditional intensity function.

Another idea supposes that the probability of spiking is determined only by the current experimental time and the time since the last spiking occurrence. This idea led to another class of probability models, proposed by Kass and Ventura (2001), the inhomogeneous Markov interval (IMI) process. The simplification proposed is to have the conditional intensities take the form

$$\lambda(t|s(t)) = \lambda(t, t - s_*(t))$$

where $s_*(t)$ is the last spike preceding time t.

2.3.1 Multiplicitive IMI Model

A subclass of the IMI processes that may be treated nonparametrically are the multiplicative IMI processes, here expressed in the additive form of the model:

$$\log\lambda(t, t - s_*(t)) = \log\lambda_1(t) + \log\lambda_2(t - s_*(t)).$$

The first factor describes the firing rate only as a function of the experimental clock time while the second factor represents the non-Poisson spiking behavior. Note that it is only dependent upon the time since the last spiking incident. Extensions of this model include dependencies upon spiking incidents farther back in the spiking history.

Fitting the IMI processes may be performed by using generalized linear models. The spike train may be approximated by discretizing time into small intervals of length Δt and converting the spike train into a series of zeros and ones. IMI processes may then be fit with standard methods. For their paper, Kass and Ventura fit the additive form of the multiplicative IMI model using cubic splines. The coefficients of the spline basis elements are determined using maximum likelihood.

2.3.2 Time-Rescaled Renewal Process

Another model that is a subclass of the inhomogeneous Markov interval processes takes advantage of the time-rescaling theorem. This theorem is well-known in probability theory and states that any inhomogeneous Poisson process may be rescaled or transformed into a homogenous Poisson process with unit rate (Taylor and Karlin, 1994). Meyer (1969) and Papangelou (1972) stated the general time-rescaling theorem, which expanded the theorem to state that any point process with an integrable rate function may be rescaled to a Poisson process with unit rate. In short, in $0 < s_1 < s_2 < ... < s_n < T$ is a realization from a point process with positive conditional intensity function $\lambda(t|H_t)$, the transformation

$$\Lambda(s_k) = \int_0^{s_k} \lambda(s|H_s)ds \tag{3}$$

yields a sequence of $\Lambda(s_k)$ which are a Poisson process with unit rate (Brown, et al. 2001).

Models based on time-rescaled renewal processes have the form

$$\lambda(t|H_t) = \lambda_0(t)h_0(\Lambda_0(t) - \Lambda_0(s_*))$$

for some function $\lambda_0(t)$. Here, h_0 is the hazard function for the ISI distribution of a renewal process and $\Lambda_0(t) = \int_0^t \lambda_0(t)du$.

2.4 Oscillatory Effects

As discussed earlier, scientists have observed that certain measurements of brain activity demonstrate oscillatory effects. Figure 2 is the local field potential which has been recorded in the mushroom body of a locust, observed two to four seconds after the presentation of an odor stimulus. Cursory observation clearly shows a periodic nature to the LFP. The top figure shows the raw data that have been obtained. In the bottom figure, the signal shown here has been filtered with a Butterworth bandpass filter between 5 and 55 Hz. The resulting filtered signal clearly demonstrates the sinusoidal nature of the signal, although one can note that the amplitude is not constant, and the frequency, while nearly constant, does fluctuate as a function of time.

With data such as neuronal data, we may conceptualize both the amplitude and frequency vary as a function of time, combined with additional noise:

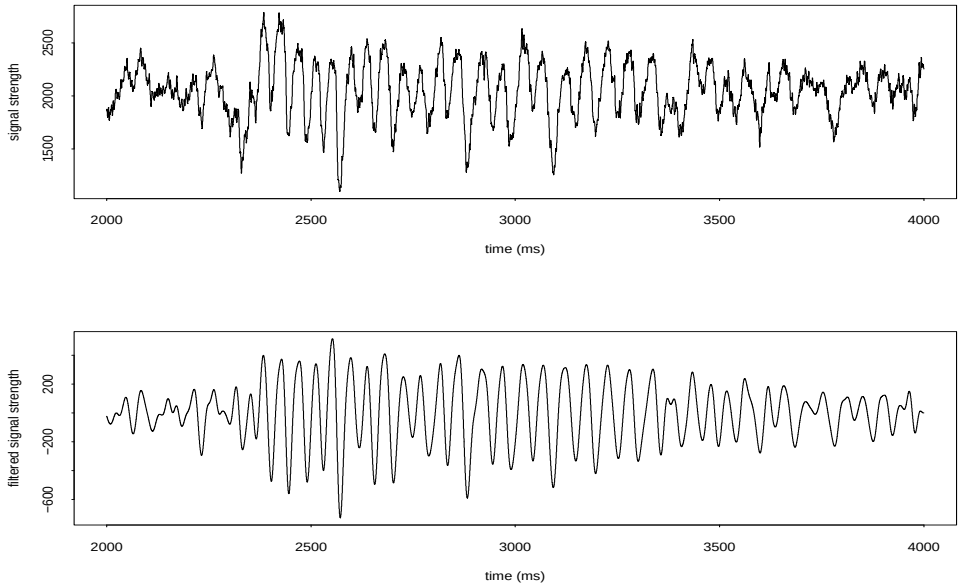

Figure 2: Local field potential in olfactory bulb of locust, recorded two to four seconds after stimulus presentation. The top figure shows the unfiltered data. The bottom figure shows the same data, after applying a bandpass filter. Both demonstrate the oscillatory behavior found in the local field potential.

$$y(t) = a(t)cos(\phi(t)) + e(t)$$

where $y(t)$ is the signal, $a(t)$ is the instantaneous amplitude, $\phi(t)$ is the instantaneous phase, and $e(t)$ is any noise present in the signal. In this case, however, there are infinite possible combinations of $a(t)$ and $\phi(t)$ that could be used to fit a signal $y(t)$. To combat this problem, several methods have been developed allowing one to extract values for $a(t)$ and $\phi(t)$. One of these methods is the Hilbert transform (Bracewell, 1965). The Hilbert transform can be considered a filter that shifts phases for all frequency components of its input, specifically shifting negative frequencies by positive 90 degrees and positive frequencies by negative 90 degrees. An analytic signal $Y(t)$ can be constructed using the Hilbert transform $h(t)$ and the original signal $y(t)$.

$$Y(t) = y(t) + i * h(t)$$

The advantage of this transform is that a unique pair of instantaneous amplitude $a(t) = \sqrt{y^2(t) + h^2(t)}$ and instantaneous phase $\phi(t) = tan^{-1}(y(t)/h(t))$ may be obtained for the signal. In addition, this transform satisfies the conditions that the amplitude and phase are continuous and are continuous in their first derivatives. Scaling of the signal by a constant does not affect the estimate of phase when the transform is applied. Also, a cosine wave with constant amplitude and frequency passed through a Hilbert transform is rendered as a sine wave with the same amplitude and frequency (Vakman, 1996). These aspects of the oscillatory behavior may then be related to the firing rates of individual neurons through generalized linear regression.

Another method used to relate spiking activity and oscillatory behavior is spike triggered averaging (Fries, et al., 2001). In this method, the oscillatory signal during a window of time surrounding each spiking event is averaged across spikes. The spectrum of the original signal is compared to the spectrum of the

spike-averaged signal, yielding the spike field coherence (SFC) at each frequency. Large values of SFC indicate a locking of the spiking to a particular phase and frequency of the oscillatory behavior.

2.5 Modeling Ensemble Effects

In addition to overall local field potentials that broadly affect neurons in the brain, individual neurons can be affected by other individual neurons. The firing of other neurons either directly or indirectly connected to that neuron may cause an increase or decrease in the firing rate. Certain neurons may be more likely to fire within a certain time interval of each other, or may inhibit their firing. Thus, a model should also incorporate ensemble effects. Understanding of how neruons interact in the brain can help to lead to better understanding of information transmittal and storage, as well as better modeling of the behavior of individual neurons.

Ventura, et al. (2005a) defined the joint firing probability of two neurons using the function $\zeta(u, v)$ where $\zeta(u, v) - 1$ is the excess proportion, above what is predicted by independence, in the probability that both neuron 1 and neuron 2 will fire in infinitesimal intervals at the respective times u and v. While this method may work for pairs of neurons, it must be extended for ensembles or groups of neurons larger than size 2.

One idea proposed to address the issues of ensemble effects was a family of loglinear models which captured interactions of all orders (Laskey and Martignon, 1996; Martignon, et al., 2000). This allows one to study the connectivity structure of a group of neurons. Expanding upon this work, Truccolo et al. (2005) formulated models for the conditional intensity function that incorporated the effects of spiking history, ensemble effects, and any extrinsic covariates x. The general framework that was used defined the conditional intensity function for a single cell j as

$$\lambda^j(t_k|H_t, x, \theta) = \lambda^j_I(t_k|H^j_t, \theta_I)\lambda^j_E(t_k|H_t, \theta_E)\lambda^j_X(t_k|x, \theta_x) \tag{4}$$

where λ_I represents the intensity function conditioned upon the spiking history of the neuron itself, λ_E represents a component related to the contribution from the surrounding ensemble of neurons, and λ_X is the component related to any extrinsic covariate.

In the framework employed in this model, the time interval $[0, T]$ is first partitioned into $j = T/\Delta t$ intervals of width Δt. The contributions from the ensemble, $\lambda^j_E(t_k|H_t, \theta_E)$ in equation 4, are then expressed in terms of a regression model of order R

$$\log(\lambda^j_E(t_k|H_t, \theta_I)) = \beta_0 + \sum_c \sum_{r=1}^{R} \beta^c_r \Delta N^c_{k-r} \tag{5}$$

where the first summation is over the ensemble of cells c which excludes the cell whose conditional intensity function is being modeled. Here, the spiking dependence of a neuron is based upon the spiking history of the other neurons in the past R time intervals.

2.6 Bayesian Adaptive Regression Splines

While producing models of the conditional intensity rates for individual neurons, it is useful to consider the method by which these rates are estimated. As demonstrated in Figure 1, the peristimulus time histogram (PSTH) is commonly used to model the time-varying rate of a point process such as a spike trian. While the PSTH successfully demonstrates the probability of firing, mentally one is able to smooth the PSTH so that one observes the smooth temporal evolution of the firing rate. This leads one to choose a method to smooth the PSTH, reducing the variability and efficiency of the PSTH. This leads one to consider curve fitting. Curve fitting is a useful technique for handling data because it allows one to produce a mathematical model to

describe the behavior of the data. In general, when applying curve fitting, one has data $(X_1, Y_1), ...(X_n, Y_n)$ where X lies on the interval $[a, b]$ and the data satisfy the following model:

$$\begin{aligned} Y_i|X_1, X_2, ...X_n &\sim p(y_i|\theta_i, \eta) \\ \theta_i &= f(X_i) \end{aligned}$$

where f is a real-valued function on $[a, b]$. Functional forms for $f(X)$ that have been used include linear regression, kernel smoothing, and smoothing splines. When choosing a smoothing technique for neuronal data, one must consider the form in which the data will appear. It is reasonable to expect the firing rate, or the intensity function, to vary slowly throughout much of the time domain of an experiment. However, it may exhibit rapid fluctuation in a relatively short interval, such as in response to a stimulus. Methods such as kernel smoothing often fail to capture these sharp changes in intensity without sacrificing smoothness elsewhere in the fit. To address this issue, certain methods have been developed, including logspline (Stone and Koo, 1986; Kooperberg and Stone, 1991, 1992; and Stone et al., 1997) and locfit (Loader, 1997).

Another particular method that was developed to address data of this form was proposed by DiMatteo, et al. (2001) called Bayesian adaptive regression splines (BARS). The method of BARS assumes the function $f(x)$ to be a spline, but the aspects of the spline are allowed to vary adaptively to the data across the domain of x. This is accomplished by allowing the knots used to fit the spline to vary according to the data. To achieve this, BARS utilizes a reversible jump markov chain Monte Carlo technique that samples from a suitable approximate posterior distribution on the knot set η. This method was first introduced by Green (1995) and generalized by Denison et al. (1998) to higher-order free knot splines. BARS uses a reversible-jump Metropolis-Hastings MCMC simulation on the pairs of k - the number of knots - and η - the location of these knots. This produces samples from the posterior on the space of splines. While BARS could be viewed simply as a procedure to find an optimal set of knots, because it generates a posterior on the space of splines it also produces an improved spline estimate based on model averaging.

3 Preliminary Results

The primary goal of this thesis pertains to modifying and evaluating existing probability models describing the firing of a neuron. This aspect has three parts - models for single neurons, models incorporating oscillatory behavior, and models for multiple neurons. Correlated with the goal of developing a model for describing the firing behavior of a neuron, the numerical methodology used to obtain estimates of the firing rate of a neuron will be addressed, namely Bayesian adaptive regression splines (BARS).

3.1 Single Neurons

When modeling the firing activity of a single neuron, the contribution of each source of variation should be included. Thus, the model for the firing rate should include components for the latency effects, excitability effects, and renewal process. When these aspects are built into the model, one notes that the firing rate of an individual neuron during a single trial k may be expressed as

$$\lambda^k(t|H_t) = \lambda(t - \tau_k)g_k(t - \tau)h(t - \tau, s_*)$$

where $\lambda(t) = E(\lambda(t|H_t))$, τ_k is the latency for the kth trial, $g_k(t)$ is the excitability component for the kth trial such that $\lambda(t)g_k(t) = E(\lambda^k(t|H_t))$, and $h(t, s_*)$ is the renewal process component for this neuron. Thus, for a sequence of spikes $s = (s_1^1, ..., s_{n_1}^1, ..., s_1^K, ..., s_{n_K}^K)$ obtained from K trials of an experiment, the probability density is given as

$$p(s|\lambda^1, ..., \lambda^K) = p(s|\tau, \lambda, \theta_g, \theta_h)$$

where $\tau = (\tau_1, ...\tau_K)$, λ is the smoothed version of the PSTH, θ_g is the set of parameters defining the excitability functions $g = (g_1, ...g_K)$, and θ_h is the set of parameters defining the renewal process function h. From this, a posterior distribution may be obtained.

$$\begin{aligned} p(\tau, \lambda, \theta_g, \theta_h|s) &\propto p(s|\tau, \lambda, \theta_g, \theta_h)\pi(\tau, \lambda, \theta_g, \theta_h) \\ &\propto p(s|\tau, \lambda, g, h, \eta)\pi(\tau)\pi(\lambda)\pi(\theta_g)\pi(\theta_h) \end{aligned}$$

where the various parameters are assumed to be a priori independent.

By drawing from the full conditionals and updating the parameter values accordingly, one can obtain a sequence of iterations that form a Markov chain with equilibrium distribution $p(\tau, \lambda, \theta_g, \theta_h|s)$. Thus, a multi-step Gibbs estimation procedure is proposed, utilizing the time-rescaling renewal process. The associated computer code has been developed, and the following outlines the algorithm used.

1. Starting values for each stage of the Gibbs sampler are determined ($m = 0$)

   (a) First, the method of Ventura (2004) is used to obtain estimates $\hat{\tau}^m = \hat{\tau}_1^m, ...\hat{\tau}_K^m$ of the latency effects for each trial $k = 1, ..K$. This method also provides an estimate of the covariance structure of the latency effects, V. These estimates are used to shift the spike trains, so that a spike from trial k at time t is shifted to time $t - \hat{\tau}_k^m$.

   (b) A smoothed estimate of the PSTH, $\hat{\lambda}_m(t|H_t)$, is obtained from the shifted spike trains using BARS and a run of about 500 to 1000 iterations. The knots and fit obtained from the final iteration are recorded.

   (c) The technique of Ventura, et al. (2005b) is used to obtain an estimate of the firing rate of the neuron for each individual trial, taking into account the excitability effects. A log linear model is fit to the data, using the smoothed estimate of the PSTH $\lambda(t|H_t)$ is used as an offset.

   $$\log\lambda^k(t|H_t) = \log\hat{\lambda}(t|H_t) + \log\bar{g}_{k_m}(t)$$

   Splines are utilized to obtain smooth estimates $\bar{g}_{k_m}(t)$. For each trial, BARS is used to fit $\log\bar{g}_{k_m}(t)$ while placing a strong prior on the number of interior knots used, limiting the number to only one or two. A run of about 500 to 1000 iterations is used. The knots and fit obtained from the final iteration are recorded.

   (d) The time-rescaled Markov model is fit, using the estimated firing rate for each trial, $\hat{\lambda}_m^k(t) = \hat{\lambda}(t - \tau_k^{\hat{m}}|H_t)\bar{g}_{k_m}(t)$ for each trial k in place of the estimate for the overall firing rate. $\hat{\lambda}_m^k(t)$ is used to rescale time for the spike times $s_1^k, ...s_{n_k}^k$ for each trial k, yielding the point process $\Lambda^k(s^k)$.

   $$\Lambda^k(s_i^k) = \int_0^{s_i} \lambda_m^k(t|H_t)dt$$

   The rescaled interspike intervals (ISIs) are obtained by differencing the resultant rescaled spkie times. BARS is used to find an estimate of the density of the rescaled ISIs, pooling the different trials. A run of about 500 to 1000 iterations is used. The results are scaled so that the resulting curve integrates to 1. This is then used to define the hazard function h_m. The knots and fit obtained from the final iteration are recorded.

2. Gibbs sampling procedure: For $m = 1, ...M$

   (a) For each trial k, a proposal latency τ_k^c is obtained by randomly generating a value from a $N(\tau_k, V_{(k,k)})$ distribution, where τ_k and $V_{(k,k)}$ are the latency and corresponding variance estimate for the kth trial obtained from the initial set-up phase. The proposed latency τ_k^c is evaluated by means of a Metropolis-Hastings algorithm. If the proposed latency is accepted, $\hat{\tau}_k^m = \tau_k^c$. If it is rejected, $\hat{\tau}_k^m = \hat{\tau}_k^{m-1}$. This obtains a draw from $p(\tau|\lambda, \theta_g, \theta_h, s)$.

   (b) A smoothed estimate of the PSTH, $\hat{\lambda}_m(t|H_t)$, is obtained shifting the spike trains by the new estimates of the latency $\hat{\tau}^m = \hat{\tau}_1^m, ..., \hat{\tau}_K^m$ and running a single iteration of BARS, with initial knots obtained from the $(m-1)$th iteration of the estimation procedure for the PSTH. This obtains a draw from $p(\lambda|\tau, \theta_g, \theta_h, s)$.

   (c) An estimate of the excitability effects $\bar{g}_{k_m}(t)$ are obtained using the smoothed estimate of the PSTH of the latency-adjusted spike trains, $\hat{\lambda}_m(t|H_t)$,c and a single iteration of BARS for each trial, using the knots saved from the $(m-1)$th iteration of the estimation procedure for the excitability functions. This obtains a draw from $p(\theta_g|\tau, \lambda, \theta_h, s)$.

   (d) The time-rescaled renewal process is fit using the estimate of the firing rate for each trial $\hat{\lambda}_m^k(t|H_t)$. This estimate is used to rescale time and a estimation of the density of the rescaled ISIs is determined using a single iteration of BARS and the knots saved from the $(m-1)$th iteration of the fit for the renewal process. This obtains a draw from $p(\theta_h|\tau, \lambda, \theta_g, s)$.

Note that to minimize identifiability issues, the Gibbs sampler is designed so that each draw of $\lambda(t)$ is drawn from a posterior distribution with mean $E(\lambda(t|H_t))$, each draw of θ_g is drawn so that the function $\lambda(t)g_k(t)$ comes from a posterior distribution with mean $E(\lambda^k(t|H_t))$, and the number of knots used to fit the splines for each function is limited to a relatively small number.

Initial results with the mechanism addressing latency, excitability, and renewal effects in the Gibbs sampler have been very promising. The technique has generated draws from the marginal posteriors of the different effects. Figure 3 illustrates a trial run of the procedure. Data were simulated from intensity functions that were shifted (latency) and scaled (excitability). The figure shows the test functions used, with the lowest intensity functions being considered the first trial, leading up to the highest intensity function as the tenth trial. The figure illustrates the estimated firing rates and credible intervals found for the PSTH and fourth and tenth trials.

The model was also applied to real data. Figure 4 shows the fit of the model to data provided by Mark Stopfer. These data were obtained from projection neurons recorded in the olfactory bulb of locusts exposed to odors, in this case hexanol. The posterior mean fit obtained for three trials of this experiment, namely the first, fifth, and fourteenth, are illustrated. In addition, error bars were added at 2500 and 3500 ms. These error bars were obtained using the 97.5 and 2.5 quantiles of the results from the Gibbs sampler. Note that while the error bars overlap at 3500 ms, the firing rate is markedly higher for the first trial at the 2500 ms mark.

3.2 Oscillatory Behavior

In addition to the trial-to-trial variability effects and the effects of the renewal process, information about existing oscillatory behavior has been incorporated into the estimation procedure. Using the Hilbert transform, one may obtain unique estimates of $a(t)$ and $\phi(t)$, the instantaneous amplitude and phase of the signal, and incorporate them into a model for the firing probability of a neuron,

$$\log(\lambda(t)) = \alpha(t) + q(\phi(t)) + r(t) * a(t). \tag{6}$$

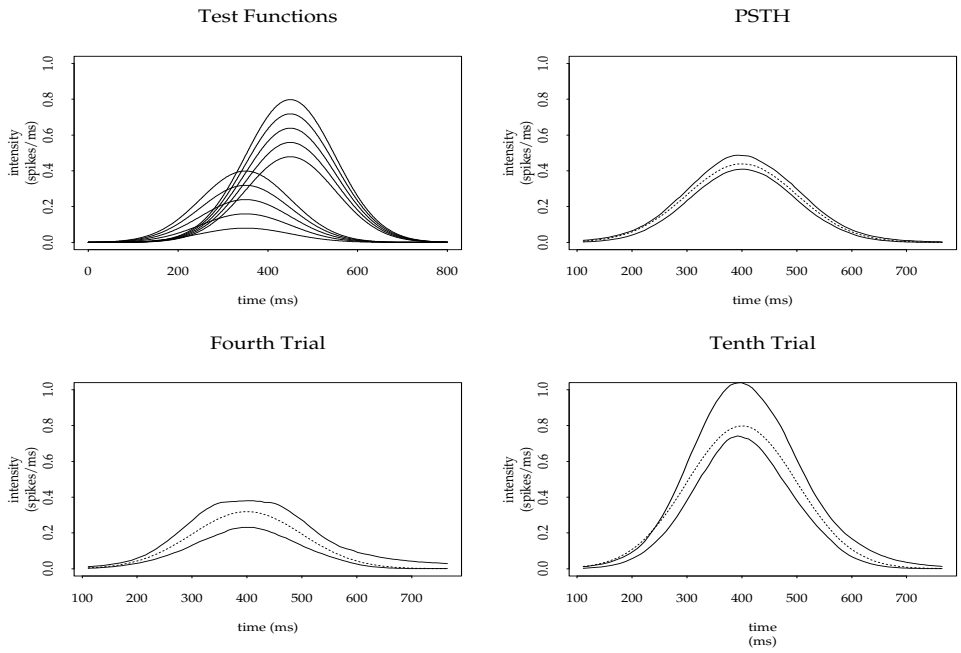

Figure 3: Early simulation results. Left top figure shows the test functions that were used, with the lowest intensity function being labeled the first trial, progressing to the highest intensity function labeled the tenth trial. The base function is a normal density with one of two latencies chosen and constant excitability. The right top figure shows the 95 percent confidence intervals for the PSTH created using the mechanism along with the true intensity function in a dashed line. The bottom two figures show the 95 confidence intervals (solid lines) created for two of the individual trials, along with the true functions (dotted lines), adjusted for latency effects.

Here, $q(t)$ is used to express the increased (or decreased) probability of firing based upon the current phase of the oscillatory behavior and $r(t)$ expresses the change in firing probability based on the amplitude of the oscillatory behavior. As before, $q(t)$ and $r(t)$ may be modeled with a spline. Generally, from preliminary results, it has been found that the amplitude of osbserved oscillatory behavior is not associated with the firing rate of individual neurons and will be omitted from future equations.

To perform the estimation technique, the Gibbs sampler described earlier is slightly modified. Here, one wishes to obtain a draw from $p(\theta_q|\tau, \lambda, \theta_g, \theta_h, s, \phi)$. Inserted after step 1d, the logarithm of the resulting firing rate for each trial $\hat{\lambda}_m^k(t|H_t)$ is used as the offset $\alpha(t)$ from equation 6, and a short run of BARS is used to find an estimate of the function $q(\phi(t))$. As before, the final fit and set of knots are saved. This is used to obtain initial values of θ_q. Inserted after step 2d, the logarithm of the fitted firing rate is again used as the offset, and a single iteration of BARS is performed in order to obtain a fit, using the knots from the previous iteration as the starting set of knots. Since current research, such as that proposed by Laurent and Stopfer (e.g. 1999) suggests that firing activity related to instantaneous phase may change as the experiment progresses, the fits for each trial are not averaged.

Initial results from data obtained from the olfactory bulb of locusts (obtained through Mark Stopfer) has shown promising connections between the instantaneous phase of the local field potential (LFP) and the firing behavior of individual neurons. Results from the earlier trials of an experiment showed little correlation between the phase of the LFP and increased firing probability. However, later trials showed a marked increase in the probability of firing slightly after the peak of the LFP. Figure 5 shows the scaling

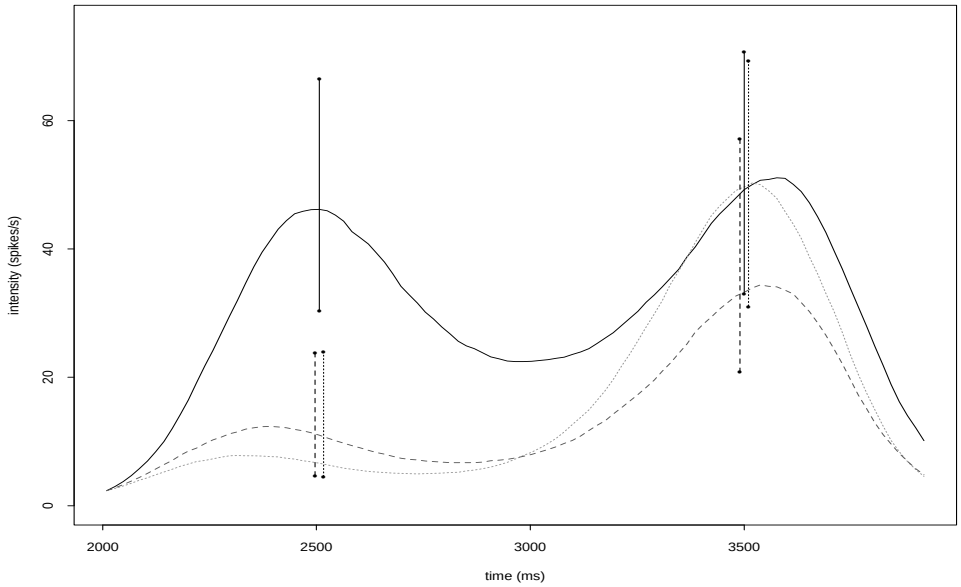

Figure 4: Fitted firing rates for three trials from a single neuron, exhibiting excitability effects. Solid line represents the first trial, dotted line the fifth trial, and dashed line the fourteenth trial in the experiment. 95 percent error bars have been added at 2500 and 3500 ms, with line type of bars corresponding to the different trials.

factor for the probability of firing obtained for four of the fifteen trials of the experiment, namely the first, fifth, tenth, and fifteenth. Note that this scaling factor is equal to $e^{q(\phi(t))}$, where $q(\phi(t))$ is as in equation 6. The estimates for the first and fifth trials (green asterisks and blue circles, respectively), show no correlation between phase and increased firing probability, and are almost exactly 1 for each value of the instantaneous phase, indicating no increased or decreased firing probability. However, the estimates for the tenth and fifteenth trials show a significant increase in firing probability approximately 30 degrees after the peak of the LFP, along with a marked decrease in firing probability approximately 90 degrees before the next peak of the LFP. This suggests that the firing of the neuron is synchronizing with the LFP as the experiment progresses. The existing estimation method does not utilize smoothing techniques specifically designed for circular data, resulting in the discontinuity at 180 degrees. This issue will be addressed in future modifications of the procedure.

3.3 BARS

One of the advantages of BARS is that it provides measures of uncertainty regarding aspects of the fit. Since BARS provides draws from the posterior of both the spline fits and functionals of the fits, credible intervals may be created for these measures. Thus, the Bayesian formalism of BARS allows for a mechanism for making inferences. However, there has been some concern voiced that the credible intervals supplied by BARS may have problems, based primarily on the absence of studies detailing their performance. Countering this argument, one may argue that the reason confidence intervals have not been created for Bayesian methods based on MCMC such as BARS is simply the computing time required to perform such a simulation at this time. For confidence intervals, to obtain a simulation standard error of 0.0025, in order to get two-digit accuracty fo an event that has approximate probability 0.95, at least 7600 samples would be required, and these studies should include cases with 100,000 MCMC iterations. However, preliminary studies suggest

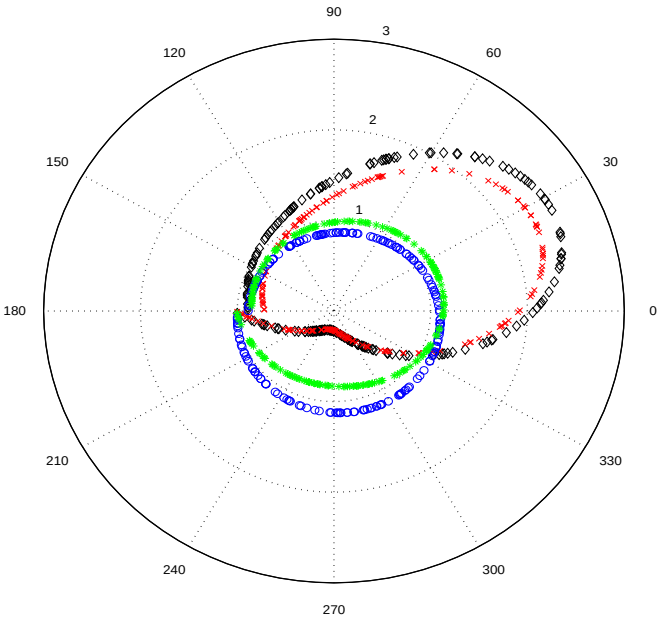

Figure 5: Association between LFP and the firing of a projection neuron, recorded in the olfactory bulb of a locust following stimulation by the odor hexanol. The rose plot illustrates the scaled increase in firing probability as a function of the instantaneous phase of the LFP, measured in degrees. 0 degrees is equivalent to the peak in the LFP. Note: green asterisks indicate the first trial, blue circles indicate the fifth trial, red 'x's indicate the tenth trial, and black diamonds indicate fifteenth trial.

that credible intervals formed through BARS have been very close to the correct frequentist percentages (Wallstrom et al. 2005).

4 Proposed Work

In a continuation of the research being performed for this thesis, various aspects for the models will be analyzed and expanded. The primary objective will be to evaluate the model that has been described. Simulation studies will be performed to assess the quality of the model. Data sets will be analyzed in the hope of finding significant effects in the estimated firing rates. The model will also be extended to incorporate multiple neurons and modifications to the renewal process. In addition, the performance of BARS will continue to be evaluated.

4.1 Single Neurons

First, simulation studies will be performed to evaluate the performance of the model. Data will be generated from various overall spiking models such as the normal distribution example shown, stepwise functions, and the simulated firing rate of a neuron. Latency structure will be varied to determine the level of latency that can be detected as significant. Excitability functions that will be used include both constant and non-constant functions. Different densities of the interspike interval will be used including non-parametric representations. During the simulation studies, the quality of the model will be assessed using measures such as mean integrated squared error (MISE) as well as coverage probabilities for the confidence bands provided by the

mechanism. In addition, goodness of fit will be evaluated, using measures such as K-S distance and plots (Conover, 1999).

After testing on simulated data has been completed satisfactorily, the model will continue to be applied to actual data sets, specifically the data supplied by Mark Stopfer from the National Institutes of Health. These data contain recordings from projection neurons recorded in the olfactory bulb of locusts which have been exposed to different odors and different concentrations of these odors. As mentioned previously, Stopfer and Laurent have observed several changes in the effects within their experiments with locusts. A significant decrease in the firing activity of neurons in the olfactory bulb following repeated exposure to stimuli was observed. This would correspond to a decrease in excitability, a result that has also been seen in the preliminary results exhibited in figure 4. Tests will be performed to determine whether or not a significant change in excitability is detected among different trials in a single experiment. In addition, changes in the firing rate across concentrations of stimulus will be analyzed, determining whether or not there is a change in firing rate.

4.2 Oscillatory Effects

As in the case of the single neuron, simulation studies will be performed to assess the quality of the model when oscillatory effects are included. Comparisons will be made to existing methodology such as that of Fries to determine whether or not a locking to a particular instantaneous phase is detectable. In addition, testing procedures will be developed to determine whether there is a preference or locking of a neuron to a particular phase of an oscillatory function. A similar test, the Rayleigh test, can be used to determine whether or not the distribution of a series of points on the unit circle is uniform across phase (Brazier, 1994). This test would require determining whether or not a continuous function is significantly greater at a particular phase or range of phases on the unit circle.

Furthermore, the Gibbs sampler incorporating oscillatory effects will continue to be applied to data obtained from Mark Stopfer. First, the development of a coherent oscillatory signal through the progression of an experiment will be evaluated, to confirm prior results. If such an osciallatory signal exists, its relationship with the firing activity of individual neurons will be examined. In initial results, examination of that data has shown that, while for initial trials there does not appear to be a significant relationship between firing of individual neurons and the phase of the LFP recorded in the mushroom body, as the experiment progresses it appears that the firing of the neuron tends to lock into a particular phase of the LFP, as illustrated in figure 5. Further study is necessary to establish the statistical significance of these observations and the prevalence of this behavior throughout the data set will be examined.

4.3 Extensions of Model

4.3.1 History Dependence

The time-renewal process may also be adjusted to account for effects that are not dependent upon only the last spiking occurrence, but upon previous spiking occurrences as well. In the multiplicative IMI model, this may be represented as

$$\log\lambda(t, t - s_*(t)) = \log\lambda_1(t) + \log\lambda_2(t - s_*(t)) + \log\lambda_3(t - s_{**})$$

where s_{**} represents the time of the spike previous to s_*, similar to the analysis performed by Kass and Ventura (2001), but now incorporating concepts of latency and excitability into the model. For the time-rescaled renewal process, the renewal process component will now be expressed as a multivariate function, such as

$$\lambda(t|H_t) = \lambda(t)h(\Lambda(t) - \Lambda(s_*), \Lambda(s_*) - \Lambda(s_{**}))$$

in the case of the extension to the previous two spiking events. Actual olfactory experiment data will be used, examining the necessity of extending the modeling process to prior spiking events beyond the last spiking occurrence.

4.3.2 Estimating Instantaneous Phase

A major piece of proposed work involves the incorporation of oscillatory data into the firing probability of a single neuron. At this stage, the instantaneous phase taken through the Hilbert transform is assumed to be given, with no inherent error. However, this is not true, since there is clearly error in the signal that is obtained, and in order to obtain the Hilbert transform a smoothing process must first be performed. Thus, two techniques are proposed to deal with this problem. The first would be to smooth the oscillatory function using a smoothing mechanism such as BARS, followed by an application of the Hilbert transform. This process would yield a draw from the posterior distribution of the instantaneous amplitude and phase at each time point t. This process can be incorporated into the Gibbs sampler when obtaining estimates of the effect of the oscillatory behavior on firing probability. The second technique involves shifting from instantaneous phase to time since last peak or trough. One might choose to use registration points, recording peaks and troughs in the oscillatory behavior, and then estimate instantaneous phase based on the time until the next peak or trough, the method originally used by the Stopfer lab, with estimates of the time of the peak of the signal obtained through a method such as BARS. This method provides the basis for a new proposal, which seeks to distance itself from the vocabulary surrounding oscillatory functions. Instead of dealing with phase, one might be interested in the time difference between the time in question and the last peak or trough. This allows one to completely avoid the problems encountered with estimating instantaneous phase. One can simply use a smoothing mechanism such as a spline to estimate the oscillatory function, and then use registration methods to determine the time to last peak. This may then be incorporated into a model for firing probability as follows:

$$\log(\lambda(t)) = \alpha(t) + p(t-u)$$

where u is the time of the last peak (or trough) and $p(t)$ is a function relating this time to the firing probability of an individual neuron. This could be done by smoothing the signal with a process like BARS, which can be used to obtain estimates of the peak location as well as information about the error of this estimate. This information can then be incorporated into the estimate of the effect on this timing between spikes and the peaks of the LFP or other oscillatory behavior with the firing behavior of individual neurons. These methods would incorporate the error inherent in estimating phase or peak time from the signal.

4.3.3 Multiple Neurons

The method of Truccolo et al. (2005) provides an excellent way of describing the temporal dependence of the spiking of a neuron upon the neurons own spiking and the spiking of the surrounding neuronal ensemble. Unfortunately, this method, requires fitting several parameters for each neuron in the ensemble. In addition, it may be beneficial in some manner to assume structure for the β parameters, which would help to reduce the number of parameters. Also, a continuous time expression would be useful. To address this, a new model is proposed for the ensemble effect:

$$\log(\lambda_E^j(t|H_t)) = \alpha(t) + \sum_c \sum_{(s<t)\in S^c} m^c(t-s) \tag{7}$$

where $m(t)$ is a smooth function of time, either a kernel function or a spline and where S^c is the set of firing times of neuron c in the set of neighboring neurons . Here, the effect of multiple firing of another neuron is

assumed to be cumulative. over the entire spiking history. The model may be modified to include only the most recent time interval, as the effect may decay over time. This can be achieved by requiring the function $m(t)$ to fade to zero as the time $t - s$ increases. In addition, the model may also be modified to look at the effect of only the last firing of either each neuron in the ensemble:

$$\log(\lambda_E^j(t|H_t)) = \alpha(t) + \sum_c m^c(t - s_*^c)$$

where s_*^c is the time of the last spike of neuron c prior to time t. As in the case of autoinhibition or autoexcitation described in section 4.3.1, this may be extended to include firing events beyond only the most recent.

This model for the ensemble effects can be estimated by modifying the Gibbs sampling technique developed for a single neuron. Here, one wishes to obtain a draw from $p(\theta_m|\tau, \lambda, \theta_g, \theta_h, s)$. After firing rates for each neuron are determined separate from the ensemble effects, they are used as the offset $\alpha(t)$ in equation 7. The functions $m^c(t - s)$ are then fit with a BARS-like iteration procedure, similar to that used to fit the PSTH and excitability effects in the model for a single neuron. A short run of BARS is used to set the initial fit and initial set of knots. Then, single iterations of BARS are used to obtain fits for the ensemble effect during the subsequent iterations of the Gibbs sampling mechanism.

Preliminary results have shown that a model incorporating ensemble effects has performed adequately. However, the performance of the model must be further assessed. Simulated data will be employed. The method will be compared to that of Brown, et al. through measures such as MISE. Although the method of Brown, et al. should perform better in certain regards, due to its lack of structure dependency in its temporal component, the methods will be compared to determine how the simpler, structured model works in comparison. In addition, this model will be applied to the locust data supplied by Mark Stopfer, examining the relationship between different neurons in the olfactory bulb of the locust.

4.4 BARS

The performance of BARS over a range of MCMC burn-in iteration and simulation iterations will be analyzed. The number of MCMC iterations will also be tested to determine how many samples are necessary to adequately sample from this posterior. To present, for a function such as the simulated firing rate of a neuron, a burn-in of 200 iterations and 2000 simulation iterations have shown to be more than sufficient to provide an adequate sampling. However, the iteration requirements for BARS will be examined for several different test functions. These include, a scaled version of the normal density, a frequency and amplitude modulated sine wave, a discontinuous step function, and the simulated firing rate of a neuron. In addition, the Donoho-Johnstone functions (1994), Doppler, Bumps, Blocks, and HeaviSine will be considered. Tests utilizing mean integrated squared error (MISE) and credible interval coverage will be used.

Furthermore, bootstrapping methods will be added to the analysis, comparing whether or not BARS without the addition of bootstrapping performs as well as with the addition of bootstrapping. It has been proposed that bootstrapping the results obtained from BARS may aid with credible intervals, and may also reduce the number of burn-in and simulation iterations required to obtain similar results. Comparisons will be made between BARS and other existing curve fitting methods, including smoothing splines and kernel regression techniques. Previous studies have focused on MISE studies comparing the performance of BARS to these established smoothing techniques. This study will replicate these results and examine confidence interval converage probability. The analysis will be refined on the Department of Statistics cluster of parallel CPUs.

References

[1] Baker, S.N. and Gerstein, G.L. (2001). Determination of response latency and its application to normalization of cross-correlation measures. *Neural Computation,* 13:1351-1377.

[2] Bear, M.F., Connors, B.W., and Paradiso, M.A. (2001). *Neuroscience: Exploring the Brain,* Baltimore: Lippincott.

[3] Bracewell, R. (1965). *The Fourier Transform and Its Applications,* New York: McGraw-Hill.

[4] Brazier, K.T.S. (1994) Confidence intervals from the Rayleigh test. *R.A.S. Monthly Notices,* 3:709.

[5] Brillinger, D.R. (1988). Maximum likelihood analysis of spike trains of interacting nerve cells. *Biological Cybernetics,* 59:189-200.

[6] Brillinger, D.R. (1992). Nerve cell spike train data analysis: A progression of technique. *Journal of the American Statistical Association,* 87:260-271.

[7] Brody, C.D. (1999). Correlations without synchrony. *Neural Computation,* 11:1537-1551.

[8] Brody, C.D. (1999). Disambiguating different covariation. *Neural Computation,* 11:1527-1535.

[9] Brown, B.N., Barbieri, R., Ventura, V., Kass, R.E., and Frank, L.M. (2002). The time-rescaling theorem and its applications to neural spike train data analysis. *Neural Computation,* 14:325-346.

[10] Conover, W.J. (1999) *Practical Nonparametric Statistics,* New York: John Wiley and Sons.

[11] Daley, D.J. and Vere-Jones, D. (1988). *An Introduction to the Theory of Point Processes,* New York: Springer-Verlag.

[12] Denison, D.G.T, Mallick, B.K., and Smith, A.F.M. (1998). Automatic Bayesian curve fitting. *Journal of the Royal Statistical Society, Volume B,* 60:330-350.

[13] Dimatteo, I., Genovese, C.R., and Kass, R.E. (2001). Bayesian curve-fitting with free-knot splines. *Biometrika,* 88:1055-1071.

[14] Fries, P., Reynolds, J.H., Rorie, A.E., and Desimone, R. Modulation of Oscillatory Neuronal Synchronization by Selective Visual Attention. *Science,* 291:1560-1563.

[15] Gabbiani, F. and Koch, C. (1998). Principles of spike train analysis. In C. Koch and G. Segev, eds., *Methods in Neuronal Modeling,* pp. 313-360. MIT Press.

[16] Gerstein, G.L. and Mandelbrot, B. (1964). Random walk models for the spike activity of a single neuron. *Biophysical Journal,* 4:41-68.

[17] Green, P.J. (1995). Reversible jump markov chain Monte Carlo computation and Bayesian model determination. *Biometrika,* 82:711-732.

[18] Hansen, M.H. and Kooperberg, C. (2002). Strategies for spline adaptation (with discussion). *Statistical Science,* 17:2-51.

[19] Hasselmo, M.E., Bodelon, C., and Wyble, B.P. (2002). A proposed function for hippocampal theta rhythm: separate phases of encoding and retrieval enhance reversal of prior learning. *Neural Computation,* 14:793-817.

[20] Kandel, E.R., Schwartz, J.H., and Jessell, T.M. (2000). *Principles of Neural Science,* McGraw-Hill Medical: New York.

[21] Kass, R.E., Ventura, V. and Cai, C. (2003). Statistical smoothing of neuronal data, *Network: Computation in Neural Systems,* 14:5-15.

[22] Kooperberg, C. and Stone, C.J. (1992). Logspline density estimation for censored data. *Journal of Computational and Graphical Statistics,* 1:301-328.

[23] Kooperberg, C. and Stone, C.J. (1991). A study of logspline density estimation. *Computational Statistics and Data Analysis,* 12:327-347.

[24] Laskey, K. and Martignon, L. (1996). Bayesian learning of loglinear models for neural connectivity. *Artificial Intelligence,* 12:213-221.

[25] Laurent, G. (2002) Oflactory network dynaics and the coding of multidimensional signal, *Nature Reviews Neuroscience,* 3:884-895.

[26] Martignon, L., Deco, G., Laskey, K., Diamond, M., Freiwald, W., and Vaadia, E. (2000). Neural coding: Higher-order temporal patterns in the neurostatistics of cell assemblies. *Neural Computation,* 12:2621-2653.

[27] Meyer, P. (1969). Démonstation simplifiée d'un théorème de Knight. In *Séminaire probabilité V* pp. 191-195. New York: Springer-Verlag.

[28] Nawrot, M., Aertsen, A., and Rotter, S. (1999). Single-trial estimation of neuronal firing rates. *Journal of Neuroscience Methods,* 94:81-91.

[29] Papangelou, F. (1972) Integrability of expected increments of point processes and a related random change of scale. *Transactions of the American Mathematical Society,* 165:483-506.

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

[31] Smith, A.C. and Brown, E.N. (2003). Estimating a state-space model from point process observations. *Neural Computation,* 15:965-991.

[32] Stone, C.J., Hansen, M. Kooperberg, C. and Truong, Y.K. (1997). Polynomial splines and their tensor products in extended linear modeling (with discussion). *The Annals of Statistics,* 25:1371-1470.

[33] Stone, C.J. and Koo, C.-Y. (1986). Logspline density estimation. *AMS Contemporary Mathematics Series,* 59:1-15.

[34] Stopfer, M., Jayaraman, V. and Laurent, G. (2003). Intensity versus identity coding in an olfactory system. *Neuron,* 39: 991-1004.

[35] Stopfer, M. and Laurent, G. (1999). Short-term memory in olfactory network dynamics. *Nature,* 402:664-668.

[36] Taylor, H.M. and Karlin, S. (1994). *An introduction to stochastic modeling* (rev. ed.). San Diego, CA: Academic Press.

[37] Vakman, D. (1996). On the analytic signal, the Teager-Kaiser energy algorithm, and other methods for defining amplitude and frequency. *IEEE Transactions on Signal Processing,* 44:791-797.

[38] Ventura, V. (2004). Testing for and estimating latency effects for Poisson and non-Poisson spike trains. *Neural Computation,* 16:2323-2350.

[39] Ventura, V., Cai, C., and Kass, R.E. (2005). Statistical assessment of time-varying dependence between two neurons. *Journal of Neurophysiology,* 94:2940-2947.

[40] Ventura, V., Cai, C., and Kass, R.E. (2005). Trial-to-trial variability and its effect on time-varying dependence between two neurons. *Journal of Neurophysiology,* 94:2928-2939.

[41] Wallstrom, G., Liebner, J. and Kass, R.E. (2005). An implementation of Bayesian Adaptive Regression Splines (BARS) with S and R wrappers, under revision for Journal of Statistical Software.

[42] Woody, C.D. (1967). Characterization of an adaptive filter for the analysis of variable latency neuroelectric signals. *Medical and Biological Engineering and Computing,* 5:539-553.