Analysis of Temporal Structure in Spike Trains of Visual Cortical Area MT

Thesis by Wyeth Bair

In Partial Fulfillment of the Requirements

for the Degree of

Doctor of Philosophy

California Institute of Technology

Pasadena, California

1996

(Defended June 9, 1995)

Acknowledgements

I wish that I could be strong enough to do things on grounds of principle alone, but, in fact, I find that some challenges are nearly impossible unless there is somebody whom I can respect and put faith in. Christof Koch was such a person that I would have followed him down many paths, but I cannot say how lucky or proud I am to have studied under him to become a scientist and to think about the brain. An equally important influence was Prof. William T. Newsome. After having met him, I could no longer espouse a purely theoretical, abstract, or detached approach to studying the brain. Nothing could have been more valuable than spending time in his lab and having access to his experience with and intuition about area MT. With their strength, integrity, discipline, and energy, I know that tremendous problems will be solved. I only hope that some of their qualities will be reflected in my future work. I can have no greater measure of my own success than my ability to provide for others what they have provided for me.

Any knowledge that has been gained here is a direct result of the data that was collected in the Newsome lab at Stanford. Kenneth Britten, Gregory Horwitz, Eyal Seidemann, Michael Shadlen, and Ehud Zohary contributed greatly in this respect. Kenneth Britten, Michael Shadlen, and Ehud Zohary have contributed immeasurably

through many conversations.

Thanks to the members of the Koch lab, especially Timothy Horiuchi, Frank Perez and John Harris, for making the lab a better place to live! I owe great thanks to Grace Chang for critically reading this thesis, and to my thesis committee for their helpful comments. I am grateful to have had the close friendship of John LeMoncheck during my stay at Caltech. Most of all, I am grateful to Stacey Fonas for her friendship and love.

This work was supported by the Office of Naval Research, the Air Force Office of Scientific Research, and the National Science Foundation. W. B. received support from a National Science Foundation Graduate Fellowship and also from the Lawrence Hanson Foundation. I would particularly like to thank Lawrence Hanson for his support, and for the stimulating conversation at the Athenaeum.

If my father has ever wondered why I came to him so few times for advice during my graduate studies, there are two reasons. For one, I had Professor Joel N. Franklin providing fatherly advice so intuitive that it typically presaged the problem for which it was needed. But also, as I gradually became aware, my father had told me almost everything I needed to know before I ever left home. From my father is derived my scientific, logical side and the drive which I have needed to complete this work. My mother, however, has forged my artistic, critical, and passionate side. Her side was always called upon when the logical side was spent, and she will always be my greatest source of strength.

Abstract

The temporal structure of neuronal spike trains in the visual cortex can provide detailed information about the stimulus and about the neuronal implementation of visual processing. Spike trains recorded from the macaque motion area MT in previous studies (Newsome et al., 1989a; Britten et al., 1992; Zohary et al., 1994) are analyzed here in the context of the dynamic random dot stimulus which was used to evoke them. If the stimulus is incoherent, the spike trains can be highly modulated and precisely locked in time to the stimulus. In contrast, the coherent motion stimulus creates little or no temporal modulation and allows us to study patterns in the spike train that may be intrinsic to the cortical circuitry in area MT. Long gaps in the spike train evoked by the preferred direction motion stimulus are found, and they appear to be symmetrical to bursts in the response to the anti-preferred direction of motion. A novel cross-correlation technique is used to establish that the gaps are correlated between pairs of neurons. Temporal modulation is also found in psychophysical experiments using a modified stimulus. A model is made that can account for the temporal modulation in terms of the computational theory of biological image motion processing. A frequency domain analysis of the stimulus reveals that it contains a repeated power spectrum that may account for psychophysical and

electrophysiological observations.

Some neurons tend to fire bursts of action potentials while others avoid burst firing. Using numerical and analytical models of spike trains as Poisson processes with the addition of refractory periods and bursting, we are able to account for peaks in the power spectrum near 40 Hz without assuming the existence of an underlying oscillatory signal. A preliminary examination of the local field potential reveals that stimulus-locked oscillation appears briefly at the beginning of the trial.

Contents

List of Figures

Chapter 1

Introduction

Our brains process information transmitted from our senses in the form of brief electrical impulses (action potentials), typically less than one millisecond in duration, that are propagated along fibers (axons). In vision, audition, and somatic sensation, the trains of action potentials that arrive from the sensory periphery are imprinted with the temporal pattern of the stimulus following its interaction with the sensory apparatus. For example, as you drive a car, fixate in the distance, and listen to the radio, your optic nerve carries a volley of action potentials each time a white dashed line crosses your retina, your auditory nerve impulses are locked in rhythm with the music on the radio, and the peripheral nerves in your spine relay the vibrations of the car to the thalamus and then on to the cerebral cortex.

It is not generally known to what extent these patterns are preserved or modified in the cerebral cortex. Nor is it known to what temporal resolution an action potential must be precise to fulfill its function. If we could add a delay randomly chosen between zero and five milliseconds to every action potential leaving the spike generating zone on the axon of all cells in your visual cortex, how much would your vision degrade? Whether the visual cortex has intrinsic rhythms, as a computer's

processing is locked to a clock, whether information is elaborately encoded for processing and transmission, as in Morse code, or whether the neuronal activity is best conceptualized as arising from a dynamical system is also not generally known.

This thesis attempts to address some of these fundamental issues by studying temporal patterns in spike trains (sequences of action potentials) that were recorded from single neurons of the middle temporal area (area MT or V5). Area MT lies in the prestriate visual cortex of the macaque monkey, roughly seven synapses away from the photoreceptors in the retina. We looked for evidence of rhythmic firing patterns and examined how precisely the spike trains were locked to fluctuations in the visual stimulus. We also compared the patterns that occurred simultaneously in spike trains of pairs of nearby neurons and measured population signals, called local field potentials, that are believed to reflect the activity of many nearby neurons. Our major findings are summarized in this thesis.

1.1 Mean Rate and Temporal Structure, a Motivation

Because the mean firing frequency in response to a sensory stimulus is reproducible under identical stimulus conditions and varies predictably and smoothly with such stimulus parameters as velocity, contrast, orientation, etc., it is commonly assumed to be the primary variable relating neuronal response to sensory experience (Adrian, 1928; Lettvin et al., 1959, or the fifth dogma of Barlow, 1972). Therefore, most studies of visual cortical neurons are not concerned with the temporal structure of spike trains and describe the output of the neuron using "tuning curves," which summarize the mean firing rate as a function of a stimulus parameter and show a peak at the preferred parameter value. This practice is supported by the existence

of a quantitative relationship between the mean firing rates of single cortical neurons and psychophysical judgements made by behaving monkeys (Werner and Mountcastle, 1963; Parker and Hawken, 1985; Barlow et al., 1987; Vogels and Orban, 1990; Zohary et al., 1990; Newsome et al., 1989a; Britten et al., 1992).

However, some electrophysiologists have focussed on the idea that the detailed dynamics of the neuronal response may carry significant information (e.g. Poggio and Viernstein, 1964; Chung et al., 1970; Strehler and Lestienne, 1986; Optican and Richmond, 1987; Abeles, 1990; Eskandar et al., 1992; Zipser et al., 1993; Bialek et al., 1991). At the time our work began, a great deal of attention had been given to the reports of stimulus-induced semi-synchronous neuronal oscillations in the 30–70 Hz range in the visual cortex of the anesthetized cat (Eckhorn et al., 1988; Gray and Singer, 1989; Gray et al., 1989; Ghose and Freeman, 1992) and the awake monkey (Kreiter and Singer, 1992; for a review see Singer, 1994).

Motivated by these findings and proposals, we set out to analyze how temporal structure in cortical spike trains could encode information regarding the stimulus or behavior. An additional motivation for studying temporal modulation developed during the course of this work. This was the idea that the detailed temporal structure of spike trains could be at least as intimately linked to the cortical architecture as tuning curves had been. Functional regions of the cortex, and even layers, have been distinguished based on the mean rate tuning of neurons. It is a powerful abstraction. However, visual neuroscience has a history of assuming that a region of the brain is homogeneous when there are not yet techniques to distinguish its sub-regions (Zeki, 1993). Therefore, we have undertaken this study of temporal structure of spike trains in the hope that it will provide a probe into the dynamics of cortical processing that will aid the fine dissecting of functional regions of the cortex, particularly at the level of individual neurons and the local cortical circuitry.

1.2 A Brief History of Our Approach

Our investigation of temporal structure relied almost exclusively on data that was collected in an ongoing series of experiments linking the responses of neurons in extrastriate area MT to the psychophysical performance of trained monkeys (Newsome et al. 1989a,b; Britten et al., 1992; Zohary et al., 1994). This database was ideal because the spike trains could be linked to both the stimulus and the response of the animal. Details of the original studies are presented in Chapter 2.

At the outset, we wanted to study the time structure of the spike trains with an eye toward the presence of 40 Hz oscillation, and therefore we relied on the power spectrum and the interspike interval (ISI) statistics. In these early analyses, we conceptualized the spike trains as somewhat stationary responses in which the mean rate was roughly constant over time, except for adaptation, and was determined by the strength of the stimulus (explained in Chapter 2). At one point, however, we made an observation that greatly changed our conceptualization of the data. We realized that the responses of many of the neurons were coupled to the dynamics of the stimulus at a time scale on the order of milliseconds. Maybe this was not initially obvious because the stimulus, by design, was dominated by white noise over a large range of parameters. But maybe, also, we had been lulled by the appealing concept that the neuron magically gleaned from the network the stimulus parameter setting and then set its mean firing rate according to its tuning curve for that parameter, but with the prescribed, well-documented, amount of noise for visual cortex.

We gradually learned that before going hunting for temporal patterns in spike trains that might result from an elaborate coding scheme, it would be prudent to first understand the output of the neurons in as much detail as possible based on their purported function. This eventually lead us to construct a simple computational model based on motion energy (Adelson and Bergen, 1985) that would help

us interpret spike trains in area MT, an area specialized for the analysis of visual motion, as indicated by its preponderance of directionally selective neurons (Zeki, 1974; Maunsell and Van Essen, 1983) and by the motion-specific effects of lesions and electrical microstimulation (Newsome and Pare, 1988; Salzman et al., 1992).

However, before constructing the computational model, we attempted to address the cause of the temporal modulation using human psychophysical experiments. It seemed likely that fluctuations in the spike trains of MT neurons would be reflected in fluctuations in the monkey's psychophysical response to the motion signal provided by the stimulus because of the link established in the original studies between the psychophysical performance and the neuronal spike count. Using modified stimuli in human psychophysics, we were able to verify that the stimulus induced measurable changes in the psychophysical response over time, but it was not clear from these preliminary psychophysical results what factor in the stimulus most contributed to these changes.

Finally, when we constructed the lineaments of a computational model of motion processing, the model necessitated a numerical description of the stimulus which, because of the fast modulation in the MT spike trains, had to be accurate down to the millisecond. The model was able to account for the presence of stimulus-locked temporal structure in terms of fluctuations in the stimulus, and it lead us to a better understanding of those fluctuations. A Fourier domain analysis of the stimulus, spurred by results from the model, provides insights that are consistent with salient observations from the psychophysics and analysis of the temporal modulation. These insights need to be studied further.

From this brief account, it should be obvious that our approach was exploratory in nature, and this was largely because of the uncertainties regarding the presence and significance of temporal structure in spike trains. In total, this thesis provides a *foundation* on which we can conduct further studies of the temporal properties

of neuronal spike trains, particularly within area MT. It represents a modest step in examining the details of neuronal spike trains and an attempt to relate them to both computational theory and biophysical properties of neurons. Much of our analysis was descriptive and does not yet directly tie to either. The results should be interpreted with the understanding that this analysis was for the most part limited to data collected in previous experiments; therefore, we had little opportunity to test the implications of our findings through further experimentation. A list of the specific results and the chapters in which they appear are given in the outline at the end of this chapter.

Finally, part of designing an elegant experiment is doing so in a way that precludes unforeseen, or at least unquantified, effects from changing the results of that experiment in the near future. A reader that is familiar with the original and ongoing studies based on this data (Newsome et al, 1989a; Britten et al., 1992; Zohary et al., 1994) may wonder how the results here bear on those studies. We believe that the results of our analysis of the stimulus-induced temporal modulation are consistent with all of the major conclusions of the earlier reports. Often these consistencies occur in subtle ways, and unfortunately there is not space in this document to carefully outline all of the connections between this analysis and the past studies. In fact, it is the exactness and the elegant design of the original experiments that have made the study of the temporal structure in this database both fruitful and enjoyable.

1.3 An Outline of Chapters

The following outline summarizes each chapter and indicates references for material that has been published.

- **Chapter 2** describes the original electrophysiological studies of Newsome et al. It is necessary to cover in detail the generation of the dynamic random

dot stimulus because the spike trains of the MT cells depend on its temporal dynamics. Here we point out a fundamental observation regarding the temporal precision of responses of the MT neurons which lead to much of the work in this thesis.

- **Chapter 3** covers preliminary psychophysical studies that demonstrated the presence of temporal modulation in the psychophysical responses which may serve as a correlate for the modulation in the neuronal response.

- **Chapter 4** summarizes the analysis of the temporal modulation induced in the spike trains by the stimulus. Here we show that the area MT cells often locked to features of the stimulus with a temporal resolution on the order of 2–10 msec. However, this modulation does not appear to be present for coherently moving patterns. Material from this chapter has been submitted for publication in the journal *Neural Computation* and has been published in preliminary form in Bair and Koch (1995).

- **Chapter 5** contains analysis which focuses on the lack of stimulus-locked modulation in the response to coherent motion. This case is interesting because in the absence of structure imposed by the stimulus, the spike trains reveal to us structure that may be generated locally by cortical circuitry. We also briefly examine cross-correlation between simultaneously recorded neurons here. Part of this chapter has been accepted for publication in *Advances in Neural Information Processing*, 1995.

- **Chapter 6** presents a model of integration of motion signals that tests whether the modulation reported in Chapter 4 can be made consistent with simple spatial integration over the spatio-temporally constrained responses of directionally selective V1 neurons, which are believed to contribute a major input to area

MT.

- **Chapter 7** contains results from some of our earliest analyses. We found that some neurons fire bursts of action potentials while others do not, and we linked burst firing to the observed peak near 40 Hz in the power spectrum of the spike trains. This chapter includes a signal detection theory analysis which varies how bursts are weighted by an ideal observer. This chapter has been published in *The Journal of Neuroscience* (Bair et al., 1994; see also Bair et al., 1993).

- **Chapter 8** summarizes results from an analytical model of the power spectrum of spike trains in the presence of a refractory period. We demonstrate that the Fourier power spectra of the spike trains could be accounted for from the ISI statistics under the assumption of a renewal process. This chapter has appeared, with additional derivations, in the *SIAM Journal on Applied Mathematics* (Franklin and Bair, 1995).

- **Chapter 9** contains an analysis of a small set of local field potential data collected in a pilot study. Stimulus-locked oscillatory patterns are observed following stimulus onset and lasting roughly 100 msec.

- **Chapter 10** contains our conclusions and suggests directions for future research.

Chapter 2

Experimental Methods

This thesis relies heavily upon data recorded in previous studies conducted in Prof. William Newsome's lab. Knowledge of the details of the stimulus and experimental paradigm is crucial for the reader to understand the results presented in later chapters. Therefore, this chapter is devoted to the essential details of the experimental paradigm and the visual stimulus. For a full account of the electrophysiological methods, see Britten et al., (1992). The final section of this chapter introduces a fundamental observation about the responses of the MT neurons to the dynamic dot stimulus. This observation changed the way we thought about the spike trains, and subsequently lead to a better understanding of the neuronal as well as the psychophysical response to the dynamic dots stimulus.

2.1 Original Electrophysiological Experiments

The original experiments were designed to compare the psychophysical performance of monkeys on a two-alternative forced choice direction discrimination task to the output of neurons recorded in area MT during that task. These studies took the neuronal response to be the total spike count during the two second presentation

of the visual stimulus, a dynamic random dot pattern that is described in detail below. It was found, using signal detection theory, that the spike count from a single MT neuron provided enough information to support a theoretical performance that was on average as good as the actual performance of the monkey. In Chapter 7 we make use of the feature that both psychophysical and neuronal responses were recorded simultaneously. However, much of the analysis here does not depend upon this aspect of the paradigm.

In the previous analysis (Britten et al., 1992), psychophysical data were compiled into psychometric functions relating percent correct choices to the strength of the motion signal. A psychophysical threshold, c_{system}, was measured for the animal in each experiment, where threshold was considered to be the motion coherence level that supported 82% correct performance. This psychophysical threshold characterizes the perceptual sensitivity of the monkey to the motion signals under the specific conditions of each individual experiment. Neuronal sensitivity was measured from the responses to preferred and null direction motion obtained over a range of coherence levels. Using a method based on signal detection theory, a "neurometric function" was computed that expressed the theoretical performance of an ideal observer who judges the direction of motion in the visual stimulus based only on the responses of the MT neuron being analyzed. Computed performance of the ideal observer was plotted as percent correct choices as a function of motion coherence, and "neurometric" thresholds, c_{cell}, were extracted in the same manner described for psychometric thresholds, i.e., as the c value that supported 82% correct performance. In general, MT neurons were remarkably sensitive to the motion signals in the stochastic display. For roughly half of 216 cells, the neurometric function was statistically indistinguishable from the psychometric function measured on the same set of trials. Across the entire set of experiments, information based on counting the total number of spikes correlated well with the monkeys' behavior, and the geometric mean ratio of neu-

ronal threshold to psychophysical threshold was 1.19 (Britten et al., 1992). One goal of the present study is to determine whether some temporal characteristics of MT spike trains correlate with either the monkeys' perceptual performance or the prior measurements of neuronal sensitivity based on integrated spike counts.

2.2 Experimental Paradigm

Three adult macaque monkeys were trained to report the direction of motion in a dynamic random dot display (described in detail in the next section) in which a certain fraction, c, the *motion coherence* or just *coherence*, of the dots moved coherently at a common speed in one direction, while the remaining dots were plotted randomly.

For a given block of trials, the random dot stimuli were placed within the receptive field of a single MT neuron, and the coherent motion signal was presented in either the preferred direction of the MT neuron or in the direction 180 degrees opposite to the preferred direction, i.e., the *null direction.* The direction of motion as well as the coherence were chosen randomly from trial to trial. We used the convention that $c > 0$ indicates preferred direction stimuli while $c < 0$ indicates null direction stimuli. Thus, at $c = 0$ all dots moved randomly and the stimulus had no bias *on average* in the preferred or null direction. For $c = 1$ all dots moved together in the neuron's preferred direction, and for $c = -1$, all dots moved in the opposite direction. The $c = 0$ and $c = \pm 1$ stimuli are important to our analysis, and are considered in more detail in the Chapter 6.

The coherence c was taken from a pre-specified range of coherence values that included the psychophysical threshold. In many cases, this range included both 0 and ± 1; however, it was not uncommon to exclude the $c = \pm 1$ stimulus because coherent motion was often correctly discriminated in 100% of trials both by the monkey and by the neuron, from the standpoint of the ideal observer. For a typical experiment,

at least 210 trials were performed: 15 trials at each of 6 preferred and 6 null direction motion coherence levels plus 30 trials at $c = 0$. Far more data were obtained for some experiments since additional blocks of trials were run as long as the cell remained well isolated. It was not uncommon to record from a neuron over the course of 1–2 hours.

An individual trial began with the onset of a fixation point presented on an oscilloscope 57 cm distant from the animal. After the monkey directed its gaze toward the fixation point, the random pattern appeared within the receptive field of the MT neuron for two seconds. The monkey attended to the random dot display and judged the direction of the coherent motion signal while maintaining its gaze on the fixation point. Computer monitoring ensured that eye movements were within 0.5° of the fixation point. At the end of the viewing interval, the fixation point and the random dot stimulus were extinguished, and two light emitting diodes (LEDs) appeared corresponding to the two possible directions of the coherent motion signal. The monkey indicated its decision regarding the direction of motion by making a saccadic eye movement to the appropriate LED; correct choices were rewarded with water or juice. A correct choice was defined as one that matched the direction used to generate the coherent motion signal in the stimulus. The monkey's eye movements were monitored continuously throughout the experiment using a scleral search coil system (Robinson, 1963).

Action potentials were recorded extracellularly from 216 MT neurons while monkeys performed the direction discrimination task. Thus physiological data and psychophysical data were obtained on the same trials. On each trial the physiological data consisted of a spike train (action potential occurrence times) recorded continuously during the fixation interval, the 2 sec visual stimulation interval and the inter-trial interval. Spike times were recorded with a resolution of one millisecond. The analysis here concerns mainly the data recorded during the 2 sec visual stimulus.

2.3 The Sparse Dynamic Dot Stimulus

At first pass, this description of the stimulus may seem excessive, but we would not have to consider the generation of the stimulus in such detail if cortical neurons were not so exquisitely sensitive to seemingly minute aspects of the retinal image.

The generation of the dynamic dot stimulus is described first in the case of a $c = 0$ stimulus, and then the method of adding the coherent motion signal is described afterwards. We do not give a full account of the statistics of the stimulus, but only try to point out what artifacts are present that could potentially affect the neuronal response.

The $c = 0$ dynamic dot stimulus consisted of 0.1° diameter dots plotted at a density of 16.7 dots/degree^2/sec on a large-screen CRT monitor (Hewlett-Packard 1321B or XYtron A21-63; P4 phosphor, 0.2 cd/m^2 mean luminance). The dots were plotted randomly (with a uniform distribution) in space in a 400 degree^2 area of the screen. Each dot was lighted for 150 μsec but appeared to be present for longer because of the persistence of the visual system. Only a circular aperture of the screen, optimized for the size of the MT receptive field, was exposed, so although the dot plotting was locked to a 6.67kHz clock, the appearance of dots within the *visible* stimulus aperture was effectively asynchronous. Also, because the dots were plotted at a constant rate, the stimulus was not Poisson in time. Note that the violation from Poisson statistics increases for larger apertures, reducing the variance in the number of dots during a fixed time period. At one extreme, the complete 400 degree^2 area would have zero variance, but a very small aperture would have Poisson statistics. A typical aperture size was 10° in diameter, having an area of about 79 degrees^2. Neurons requiring aperture diameters of 15° or more, revealing nearly 50% of the entire screen, were encountered infrequently in these studies.

The coherent motion signal was introduced by plotting dots with a particular Δx

and Δy, i.e., Δr and $\Delta\theta$, and Δt with respect to previously plotted dots. The time step, Δt, was fixed at 45 msec, a value near optimal for humans and monkeys (Morgan and Ward, 1980; Newsome and Pare, 1988). Because 300 dots can be plotted in a 45 msec period, it is useful to think of the stimulus as consisting of 300 dots that were continuously cycled through and replotted. A $c = 1$ stimulus would appear to the human observer as a sparse, rigid, spatial pattern that translates behind a circular aperture; however, each individual dot could be described as jumping at 22.2 Hz $(1/\Delta t)$ but with random phase relative to the other dots. Since it is known that the sum of sine waves with a fixed frequency and random phases is also a sine wave of the same frequency, it is no surprise that some MT cells were found to have a 22.2 Hz component in their response to the $c = 1$ stimulus (see Fig. 6.9).

Lower coherence stimuli, $|c| < 1$, consisted of a mixture of dots carrying a random signal, "noise dots," and dots carrying the coherent motion signal "signal dots." Such a stimulus may be conceptualized in a sequence of 45 msec epochs. In the first epoch, 300 dots are plotted at random spatial locations (only some of which appear within the stimulus aperture). During the next 45 msec, with probability $|c|$, each dot in turn is replotted at the appropriate spatial location, i.e., shifted by Δx and Δy, to yield the velocity preferred by the cell, or, with probability $1 - |c|$, is plotted at a random location within the 400 degrees^2 area. Notice that regardless of c, the first 45 msec of each stimulus is the sequential plotting of a random pattern. This is true even for $c = \pm 1$. Only in the second 45 msec epoch of the stimulus can the coherent signal appear, and the lower the coherence and the smaller the aperture size, the longer it will be before a "signal dot" occurs.

Another aspect of the stimulus is the "lifetime," or number of 45 msec epochs for which a signal dot persists. Because in each epoch the fate of each dot is independently determined based on c, the lifetimes are distributed geometrically—the probability of a dot moving for n steps in the same direction being $|c|^{(n-1)}$.

The diameter of the circular aperture (between 5° and 15°) in which the dots appeared was optimized for the receptive field of the neuron, so neurons having smaller receptive fields were stimulated with fewer total dots. This probably did not cause a noticeable change in the level of stimulation for the MT cell since it was reported that these cells saturate with only a few dots in their receptive fields (Snowden et al., 1992).

From trial to trial, not only were the direction and coherence of the stimulus randomized, but also the exact pattern of dots was controlled by a "seed" that was used to initialize the random number generator (a pseudo-random m-sequence). This seed can be held constant to cause each stimulus, at a particular c, to be identical, or it may be changed during the course of the experiment to sample from the ensemble of all possible coherence c stimuli. Thus, if for any particular c value, the same seed was used on different trials, the monkey saw on each of these trials the identical random dot pattern. The seed values, thus the identity of the stimuli, were not stored. In Chapter 4 we consider how this limits the type of analysis which can be performed on the associated neuronal responses.

Finally, there are important differences between the dot stimulus used here and those of other studies. Some studies always plot dots with apparent matches, but vary the distribution of velocities. Note that in Fig. 1 in Britten et al. (1992), the arrows attached to the noise dots should not be taken to mean that the noise dots have been intentionally replotted to create apparent motion. In the study by Snowden et al. (1992), the dot lifetime was 0.5 sec, making the stimulus appear highly coherent, approximating $c = 96\%$.

2.4 Other Experimental Paradigms

The majority of the data examined in this thesis was collected using the stimulus and paradigm explained above, but we occasionally depended on results collected in other experiments. Those experiments are briefly described below, with focus on how each differs from the standard paradigm.

Data are included here from the "split-field" experiments (see Fig. 5.6) in which the stimulus was presented in a frame-based fashion at 40 Hz (Britten and Newsome, unpublished experiments). This is in contrast to the effectively asynchronous plotting used for the experiments described above. The cases considered here correspond to those in which the right and left halves of the stimulus were matched and the stimulus was $c = \pm 1$. Our interest in these experiments is related to the 40 Hz modulation that the frame-based stimulus induced in the responses of the cells.

Data are included from studies in which pairs of neurons were recorded simultaneously from a single electrode (Zohary et al., 1994). In these studies, the stimulus was generated on a 60 Hz frame refresh device and the Δt was 50 msec (three frames). It was found that the spike counts of neurons tend to be correlated if the spatial locations of the receptive fields overlap and if the direction preferences of the neurons are similar.

Local field potential data was collected in experiments conducted in collaboration with Greg Horwitz and the author using the same 60 Hz frame-based stimulus as described for the pairs study. This method of recording is described in Chapter 7.

2.5 Temporal Structure: A Motivation

In the original work by Newsome et al., the neuronal response was defined to be the spike count during the 2 sec stimulus. Knowledge of the finer detail concerning the

organization of the action potentials in time during the spike train was not necessary, and therefore not studied. In our study of the temporal structure of these spike trains, we observed that when stimulated on multiple occasions with the identical random dot stimulus, a neuron gave a response that could be reliable on the order of milliseconds. This observation is apparent in Fig. 2.1. The left side of the figure shows the responses to 90 different $c = 0$ stimuli, while the right side shows 90 responses to the same $c = 0$ stimulus.

This observation prompted us to ask whether the variance of the stimulus affects the monkey's decision, and whether the decision of the psychophysical observer is changing over time. This issue is covered in the next chapter. It also prompted us to ask whether such temporal modulation is consistent with current models of area MT. This question is dealt with in the Modeling chapter. In Chapters 4 and 5, we analyze the character of the temporal structure in more detail.

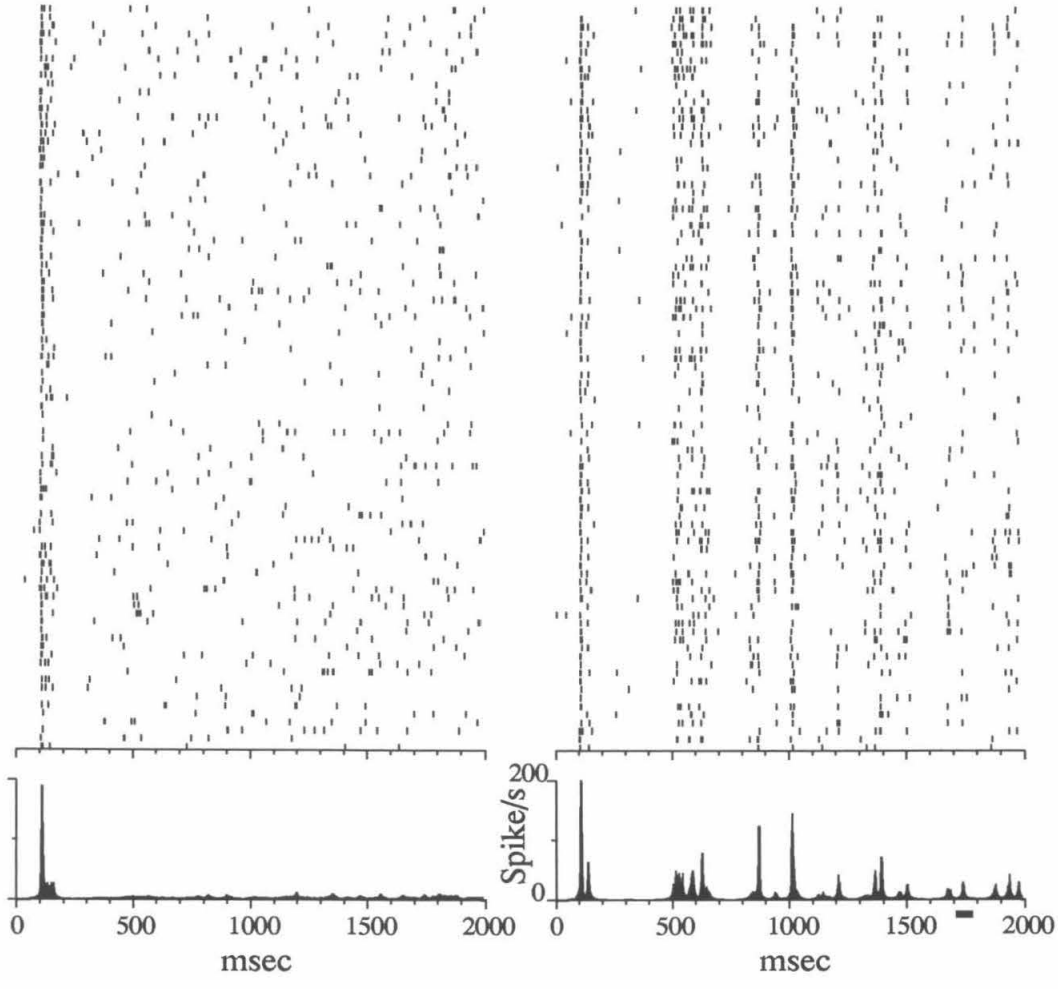

Figure 2.1: The neuronal response of one cell in area MT in a behaving macaque monkey to randomly seeded dynamic dot stimuli at $c = 0$ presented for 2 sec (left) appears to be well described by a point process with a mean rate of 3.4 Hz (excluding the initial transient). However, when a dynamic dot stimulus formed with a *particular* seed was repeated (but interleaved with different stimuli) many times, the reliability of the response becomes apparent (right). Viewing this figure from an acute angle reveals the precision of the pattern; for example, nearly all spikes in the final 400 msec of the response cluster in six vertical streaks. Below each set of spike trains is a post-stimulus time histogram (PSTH) computed from each set of 90 trials using an adaptive square window centered at each point and widened to capture 10 spikes. The set of data on the left is referred to as control data.

Chapter 3

Psychophysical Evidence for Temporal Modulation

Once the observation was made that the spike trains of the MT neurons show a marked rate modulation that was locked to the random dot stimulus, it was natural to ask whether the perception of motion also varied over time, particularly because of the match between neurometric and psychometric curves established in the original studies (Newsome et al., 1989a; Britten et al., 1992). In this chapter, we describe a few pilot psychophysical studies that were performed to verify that the psychophysical response on the two-alternative forced choice discrimination task did vary with the stimulus. These experiments are exploratory in the sense that more subjects and additional controls need to be run. The results here, however, are qualitatively consistent with the observed modulation in the MT spike trains (Chapter 4) and with the output of a motion energy model responding to the dynamic dot stimuli (Chapter 6.1). The difficulty of quantifying the psychophysically relevant motion signals in the stimulus was the primary factor that limited the scope of these early psychophysical experiments.

Originally, we believed that the trial-to-trial variance created by changing the randomization seed for the stimulus had little or no influence on the psychophysical responses and that noise of biological origin was dominant (Britten et al., 1992). However, in light of the variance induced in the neuronal response by seed changes (Fig. 2.1), we conducted further tests for variance in the psychophysical response. These results are described in the three sections that follow.

For the psychophysics described below, the dynamic dot stimulus was presented on a 60 Hz frame refresh monitor. The stimulus could be thought of as a movie because it was a sequence of frames of dots. The subjects sat 57 cm from the screen in complete darkness. The stimulus was presented at high contrast in a 5° diameter aperture centered along the horizontal axis 5° to the right of the fixation point. The speed of the signal dots was 1.67°/sec, the direction of motion was either up or down, and the stimulus duration was 2 sec. The subjects fixated a red "x" during the presentation of the stimulus and indicated their response with a keypress. A "beep" was sounded for incorrect responses, i.e., those that did not match the intended direction of motion of the stimulus. Out of habit, we will refer to a stimulus moving in particular direction as "preferred" and to a stimulus moving in the opposite direction as "null." Keep in mind that these are arbitrary designations to the psychophysical observer and have meaning only with respect to a neuron.

3.1 Human Psychophysics

A psychophysical observer of the random dot stimuli is defined by his performance (fraction correct) at each coherence level $0 \leq c \leq 1$. (We assume performance is equal for preferred and null directions of motion.) If the stimulus variance from changing the seed is drowned in the noise of the system, the response to a particular stimulus pattern on one trial will give no *additional* information about the response to that

same pattern on a later trial over and above that provided by the knowledge of c.

To test the hypothesis that two responses to a particular stimuli are uncorrelated for a particular c, we had human subjects perform the two-alternative forced choice direction discrimination paradigm (described in Chapter 2) on two separate occasions using the same set of stimuli on both occasions. Formally, let X be a random variable representing the subject's decision after viewing a stimulus of a particular coherence c, and let Y represent the subject's decision after viewing the identical stimulus at another time, possibly the next day. The particular outcomes of a given trial, x and y, take values of 0 and 1 depending whether the response was the null direction or the preferred direction, respectively. We computed the correlation coefficient (Pearson's r) between X and Y over a set of stimuli at a particular coherence level. (Note, it is trivial that the responses would be correlated if the coherence level was allowed to vary, since for large c the subject will answer correctly so that $x = y = 1$, and for c near zero, the subject will guess so that x and y are randomly assigned 1 or 0.) At a particular coherence level, the correlation coefficient, r_c, is

$$r_c = \frac{C}{\sigma_x \sigma_y} = \frac{\mathrm{E}(xy) - \mathrm{E}(x)\mathrm{E}(y)}{\sigma_x \sigma_y} \tag{3.1}$$

where C is the covariance, E indicates expected value, and σ_x and σ_y are the standard deviations of X and Y. This correlation coefficient is known as Pearson's r. It is not difficult to show that r is linear in the number of component matches, i.e., the number of times $x_i = y_i$.

Fig 3.1 (top) shows the distributions of r_c values accumulated over three subjects (whose performance had asymptoted) for c values from nearly zero to above threshold (roughly 5%). Each r_c was computed from the pairs of x and y values for 15 different stimuli at the specified c. These data were extracted from two blocks of typically 300 trials (sometimes 360) where both blocks consisted of the same set of stimuli, but

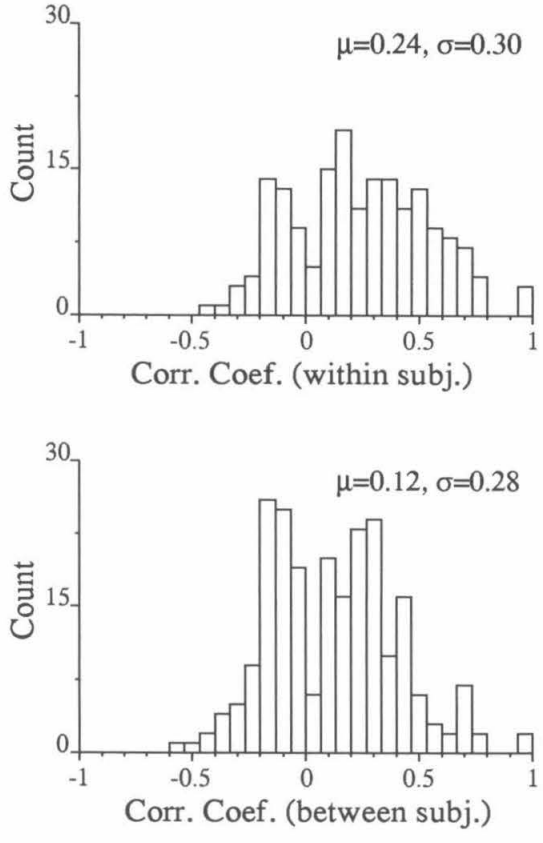

Figure 3.1: Correlation coefficients for the relationship between motion discriminations made by subjects on identical motion stimuli. **Top:** The histogram shows the distribution of r values between the subjects' decision at one time and then at a later time (typically the next day) on sets of 15 different stimuli at the same direction and motion coherence. The individual subjects' means were 0.26 (WN), 0.21 (KB), and 0.22 (WB). **Bottom:** The same is shown for correlation between the decisions of two different subjects for the same stimulus patterns. The mean r value was half as much for between-subject correlations. The dips near zero are due to the limited possible values of r_c with 15 trials.

ordered randomly. The average r_c was 0.24 (SD 0.30). The significant ($p < 0.000001$, t-test) rightward shift in the distribution of r_c away from the zero value expected under the null hypothesis was present at all coherence levels tested, as shown in Fig. 3.2. The distributions of r_c for the three subjects individually had significant ($p < 0.001$) rightward shifts, but the mean shifts (given in the caption of Fig. 3.1) were not significantly different between any pair of subjects. This result demonstrates that some stimuli at a particular coherence c give stronger impressions of motion, in either the preferred or null direction, than an average stimulus at coherence c, and that this trial-to-trial variation is not lost in the biological noise.

Computing the correlation coefficient at high motion coherence values becomes impossible if the subject gives the same, i.e., correct, answers on all trials. Without variation in the response due to mistakes, r_c is undefined. However, as long as a few mistakes are made, the correlation coefficient will remain a valid measure of whether those mistakes are induced by a particular stimuli or whether they occur at random across the ensemble of stimuli, in particular, under the null hypothesis that mistakes are uncorrelated, r_c still has an expected value of zero.

If the observed within-subject correlation for particular motion stimuli is the result of excessive motion signals in the stimulus, we would expect there to also be a *between*-subject correlation when different observers viewed the same (but randomly interleaved) sets of stimuli. Fig 3.1 (bottom) shows that there is a significant ($p < 0.000003$, t-test) but smaller r_c (mean 0.12, SD 0.28) between subjects. The between-subject correlation may be less because the subjects have consistent but different guessing strategies or simply because they process motion differently.

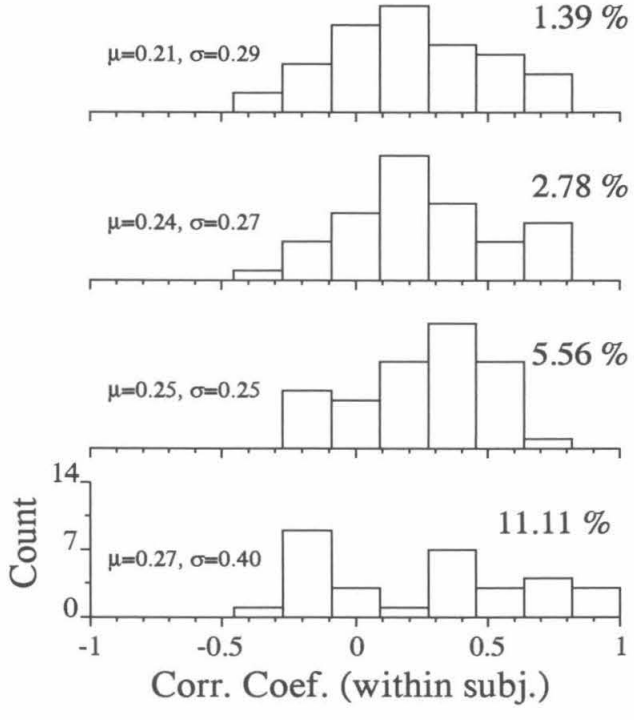

Figure 3.2: Within-subject decision correlation shown by stimulus coherence. There is no strong variation with c in the mean correlation coefficient relating a subject's decision on two presentations of the identical stimulus (see text for details). The artifactual peaks at 11.1% result from the small number of r values which are possible from 15 trials when the subjects make few (typically 1 or 2) mistakes. The coherence levels here span the range from chance to above threshold levels.

3.2 Monkey Psychophysics: A Look Back

Once it was clear that the stimulus variance affected human psychophysical decisions, we went back to analyze the monkey data. As mentioned in Chapter 2, in recordings from 25 cells in monkey E, $c = 0$ stimuli having a particular seed were interleaved with other trials. Data were also collected from trials at $c = 0$ across varying seeds. We are interested only in the psychophysical responses here; the neuronal responses will be analyzed in Chapter 4.

For the set of trials at $c = 0$ in which the seed varied, we computed the probability P_{var} that the monkey chose the preferred direction. If the monkey has no bias to guess a particular direction, P_{var} should be 0.5. We computed a similar measure P_{novar} for the $c = 0$ trials on which the stimulus seed was not varied. If an individual stimulus has substantially more preferred or null direction signal due to its random generation, then we expect P_{novar} to differ significantly from P_{var}. Fig. 3.3 (top) shows with vertical lines the change in performance from P_{var} (filled circles) to P_{novar} (open squares). The abscissa of each data pair is the statistical significance of the difference between the two probabilities computed using the G-test for independence, a log-likelihood ratio test. Six of 25 cases are more significant than 0.05. High significance values are driven jointly by a large number of trials and a large difference, $P_{var} - P_{novar}$. For the six most significant (left-most) pairs on the plot, P_{var} is near 0.5 as expected, while P_{novar} is pulled away from 0.5. For one of the 25 patterns, the $c = 0$ stimulus caused the monkey to guess the preferred direction on 93% (28 out of 30) of trials, which is greater than the 82% performance achieved by a threshold level preferred stimulus.

We tested the results of the statistical G-test using a simple Monte Carlo simulation. The bottom panel of the Fig. 3.3 shows pairs of data points generated under the null hypothesis that P_{var} and P_{novar} had a similar mean (the average of the measured

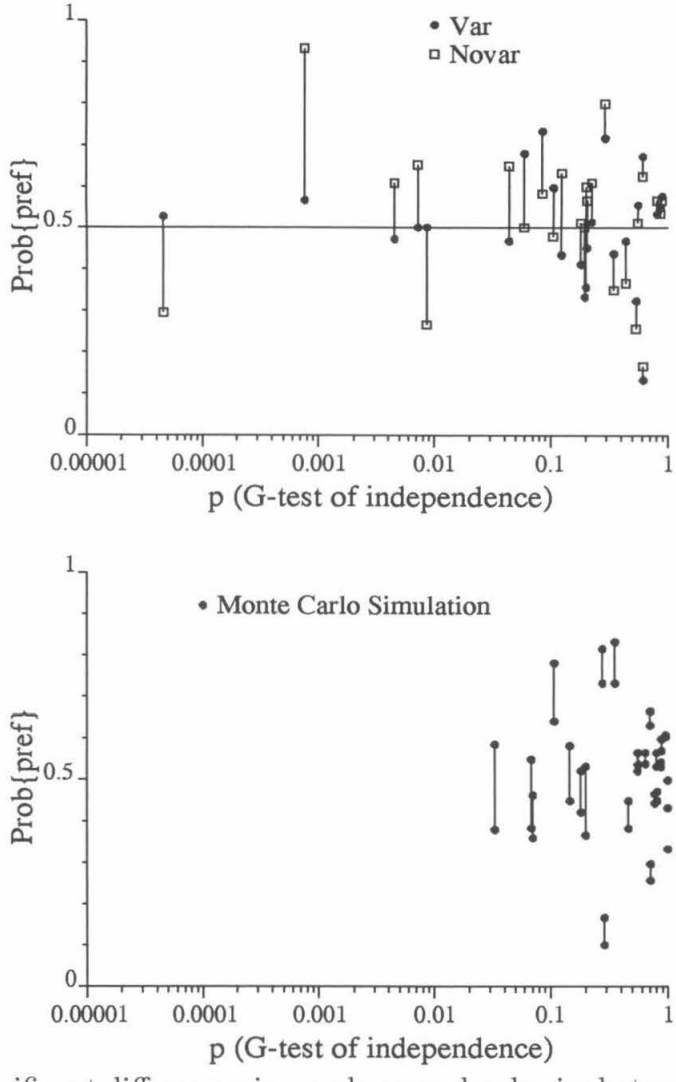

Figure 3.3: Significant differences in monkey psychophysics between the probability of preferred response across the ensemble (with respect to randomization seeds) of $c = 0$ stimuli, which is expected to be 0.5, and the probability of preferred response for a particular $c = 0$ stimulus pattern. **Top:** The difference between P_{var} (filled circles) and P_{novar} (open squares) is shown as a vertical line plotted at the significance value associated with this difference based on the G-test for independence. Six of 25 cases are significant ($p < 0.05$). **Bottom:** A Monte Carlo simulation shows that the statistical test appears to be accurate and does not find highly significant differences between two arbitrarily divided samples.

individual means). Effectively, an identical, biased coin was flipped twice and the two data points are shown connected with vertical lines. The G-test was applied to this data, and the abscissa indicates that only one point in 25 fell above 0.05 significance by chance.

Conceptually, the ensemble of $c = 0$ stimuli may be distributed along an axis which measures preferred direction motion. We expect this distribution to be centered at zero, having no bias to the preferred or null direction, and we only expect the members of the ensembles at the tails of the distribution to bias the psychophysical decision significantly. This agrees qualitatively with what we observe here and is consistent with the results from the human psychophysics.

3.3 The Cut Experiments

Returning to human psychophysics, the question of trial-to-trial variance induced by changing the stimulus seed is considered in terms of variations *during* a single trial. Does a two second stimulus at coherence level c provide a relatively constant, but possibly weak, signal that is corrupted by noise in the biological system as it is integrated to yield a final decision for the direction discrimination task? If so, at low, i.e., subthreshold, c, the observer would be guessing—performing at 50%—if forced to respond after viewing only a short portion of the stimulus. Alternatively, the stimulus might contain signals that are registered with great precision, relatively uncorrupted by the noise of the biological system. An extreme, non-biological example of this alternative is a digital computer. It never guesses; given the same set of points defining a dynamic dot stimulus, it will always produce the same answer, right or wrong. That answer may well be wrong if the digitization of the stimulus is a manifestation of noise, but for no particular stimulus (a string of bits) can its performance be 0.5.

To test for temporal modulation in the psychophysical decision, we devised the

"cut" experiments. In these experiments, only a few different stimuli were used, but the presentation of the stimuli were cut short at one of 7 predetermined locations, at which time the subject responded. For these blocks, 12 stimuli with different seeds were used, 6 (3 upward and 3 downward) at each of two (5.6% and 11.1%) coherence levels. Each stimulus had 7 different cut versions which were repeated five times each. Thus one block had 12*7*5 = 420 stimulus presentations, randomly shuffled.

Fig 3.4 shows the performance of one human subject (the author) from 4 blocks of cut experiments performed on the same day. The top shows the performance over time, i.e. at each of the seven cut times, for 3 stimulus movies at 5.6% coherence. Two striking features are the lack of convergence to correct performance over time and the occurrence of large fluctuations in decision probability with time during the movies. For the stimulus plotted with the thick line and open squares, the subject answered correctly on 18 of 20 stimulus presentations when the movie stopped at 0.5 sec (30 frames), but the subject answered incorrectly on 18 of 20 times at the 1 sec cut. Later, at the 1.75 sec cut, the subject answered 20 of 20 correctly. Apparently, the stimulus appeared to be moving upward initially, and later, at 1 sec, moved downward even though it was generated at 5.6% *upward* coherence. The performance on two other stimuli are shown in Fig 3.4 (top). For one, performance was consistently high throughout the stimulus, while the other was marked by a progressive decline in performance. These fluctuations in decision cannot simply be caused by noise because the results from the four individual blocks are consistent. Fig 3.4 (bottom) shows the individual block results that contributed to the thick line in the top part of the figure. (The lines in the bottom panel are shifted for clarity, each line represents 5 repetitions).

The psychophysical subject in Fig. 3.4 is below threshold performance at $c = 5.6\%$, achieving only 70% performance for the full 2 sec stimuli. The results from 11.1% coherence stimuli (not shown) indicate faster convergence to perfect performance, but

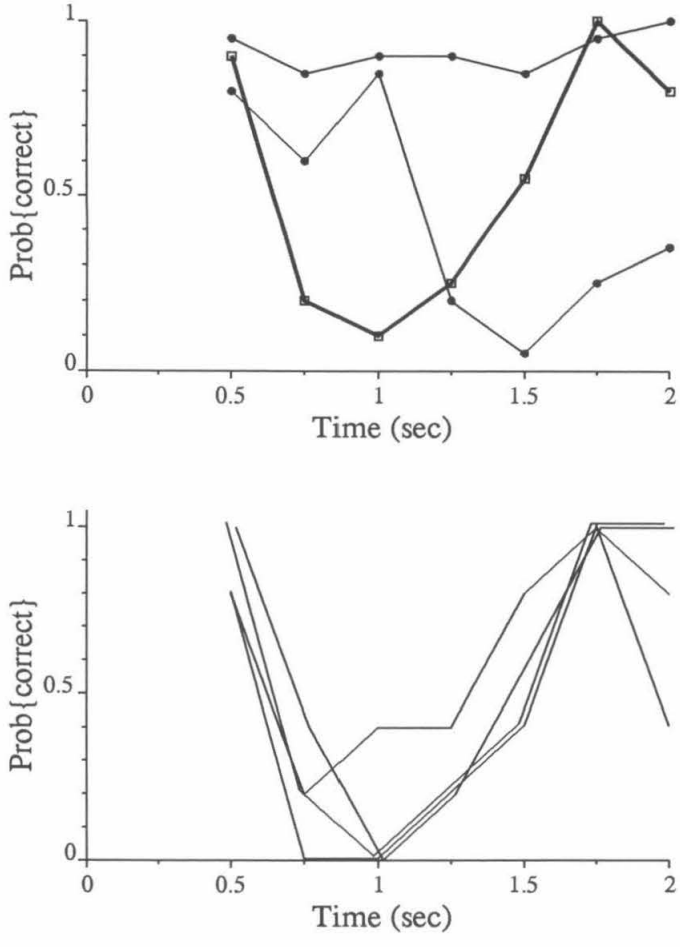

Figure 3.4: Results from "cut" experiments for $c = 5.6\%$ stimuli. **Top:** Psychophysical performance varies consistently over time for three different stimuli. The heavy line shows that after viewing a particular stimulus for 0.5 sec the subject gives the correct answer on 18 of 20 trials, but at $t = 1$ sec, the subject chooses the preferred direction on 2 of 20 trials. Later, at 1.75 sec, the subject decides correctly on 20 of 20 trials. For a second stimulus generated from a different random seed, the performance remains relatively constant at above 0.80 during the entire trial. A third stimulus produces a wide swing in performance during the trial, dropping to 1 out of 20 correct at 1.5 sec into the stimulus. **Bottom:** Results from the 4 separate blocks that make up the heavy line in the top plot. Notice the consistency from one block to another. Each block contained 5 trials at each of 7 different cut locations during the stimulus movie. The subject received auditory feedback for incorrect responses.

even at this above threshold level, particular stimuli can induce large fluctuations in performance or can result in a progressive decline in performance over 2 sec.

The stimulus in the cut experiments was centered along the horizontal axis. By inverting the stimulus about this axis, the direction of all relevant motion signals in the stimulus are reversed, but all parts of the stimulus pattern remain at the same eccentricity from the fixation point. This stimulus inversion results in an inversion of psychophysical response (thus an identical response in terms of probability correct) as expected if a motion signal is driving modulation in the response. However, if the stimulus is mirrored about the vertical axis, thus swapping far for near eccentricities, the response pattern changes. This may be in part due to the known variations in spatio-temporal sensitivity with eccentricity.

3.4 Summary and Future Work

The preliminary psychophysical observations here suggest that the noise dots (which far outnumber the signal dots for c near and below threshold) carry signals that can supersede those carried by the signal dots alone. The noise dots do not, therefore, simply act in a way to suppress the coherent signal down to the level of biological noise in the system. This appears to be true even at coherence levels above threshold. We thought it best to gain an understanding of the stimulus and its interaction with a candidate biological motion system (see Chapter 6.1) before pursuing more rigorous psychophysical experimentation.

In the future, more subjects must be tested and some controls need to be performed. For one, the cut experiments should be performed in a block-wise fashion where each block consists of equal length stimuli to determine whether knowledge of the stimulus duration influences the decision process. Also, the cut experiments should be performed with larger aperture stimuli. In this case, one prediction is

that a larger field will provide both a more consistent coherent motion signal and a more balanced noise signal, increasing the chance of 0.5 performance early during a cut trial and decreasing the systematic fluctuations between up vs. down decisions. However, if the system is highly non-linear or if it weights the coarsest spatial scale most heavily, the same results may be obtained.

For this motion discrimination task, it is evident that an observer may not be exclusively leaning to integrate subtle motion signals over space and time, but may also be learning to suppress types of motion signals that are not associated with the signal dots in the apparent motion display (e.g., Vaina et al., 1995). The analysis of the stimulus in Chapter 6 demonstrates that the signal dots have a characteristic sinusoidal structure in the frequency domain that might be learned. If the correct response—that which provided a reward to the subject—was determined by analysis of each particular stimulus using a biologically based motion detection system, the subjects are likely to perform better. It would be interesting to test whether such a system, if tuned for each observer, would vary substantially from one person to the next. Training and testing the motion system on a task that it does naturally, i.e., one based on smooth rather than apparent motion, might better assess the underlying reliability and noise level of the neural substrate of visual motion processing.

Similar work regarding the reliability of responses in the face of noise is being conducted in the auditory system for noisy inter-aural time difference stimuli (H. S. Colburn, personal communication).

Chapter 4

Precise and Reliable Temporal Modulation

Originally we conceptualized the MT responses to the dynamic random dot stimuli as random spike trains whose mean rate was set by the coherence c used to generate the motion stimulus, and we were not aware that the trial-to-trial variation in the stimulus, caused by changing the randomization seed, was reflected in the spike trains (Britten et al., 1992; Britten et al., 1993; Softky and Koch, 1993; Bair et al., 1994).

The observation that first revealed the cells' ability to follow variations in the stimulus was a rather puzzling one. When the post-stimulus time histograms (PSTHs) for a preferred and a null direction stimulus were plotted together, as in Fig. 4.1, they shared common peaks. The algorithm that generated the random dot stimuli was such that, when two oppositely directed stimuli are generated from the same random seed, as is the case here, the signal dots start at the same points in both stimuli but move in opposite directions, while the noise dots simply occur at the same points in both stimuli. If variations in the number or spatial distribution of the signal dots were dominant in setting the instantaneous firing rate of the cell, then we would have

expected the time courses of the null and preferred PSTHs to be inverted, not similar. Thus, the conclusions drawn from the psychophysical experiments in Chapter 3 were consistent with this observation—the noise dots seemed to play a significant role in creating modulation. We will take up this line of thought again in Chapter 6 when we discuss the frequency domain analysis of the stimulus.

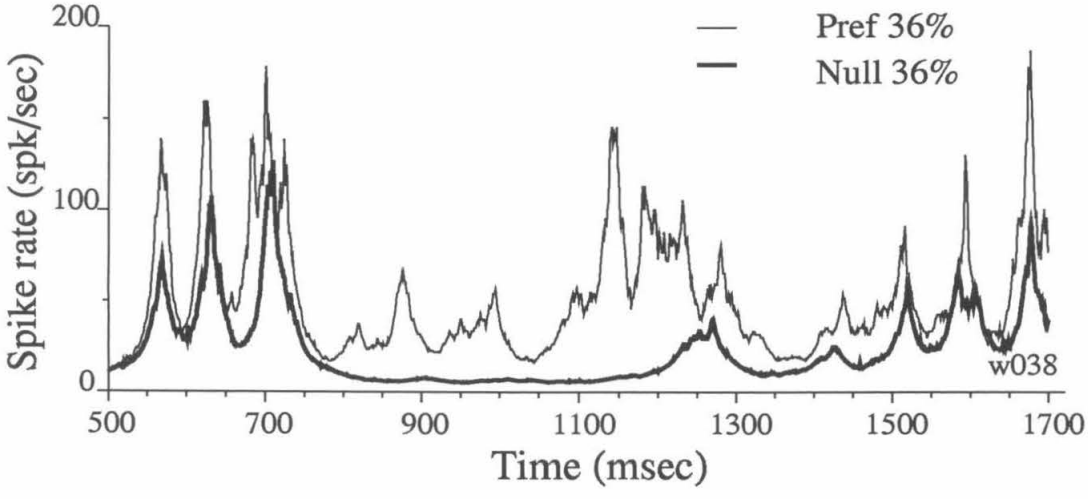

Figure 4.1: Responses of cell w038 to a particular stimulus pattern at $c = 0.36$ and $c = -0.36$. The noise dots were identical in both stimuli while the signal dots moved in opposite directions (preferred–thin line, null–thick line). If the modulated firing rate was caused by variations in the number of signal dots over time, then we would expect the modulation to be inverted for the null direction. The presence of peaks at similar times in both responses suggests that the *noise* dots create signals that strongly influence the firing rate of the MT cell. (Each PSTH was computed from 15 trials using an adaptive window to capture 10 spikes.)

This chapter will focus on quantifying the temporal modulation, without further questioning its origin. The temporal structure studied here is similar to a stimulus-locked response sometimes referred to as a "grain" response which has been reported in earlier studies of mammalian visual cortex (Tomko and Crapper, 1974; Hammond and MacKay, 1977; Gulyas et al., 1987; Snowden et al., 1992). However, the time scale and stimulus dependency of this type of modulation has not been characterized at the

trial-to-trial level. Other types of stimulus-locked modulation have been studied as well. Stimulus-locked modulation has been shown to exist in visual cortex for static patterns (Richmond et al., 1987, Richmond et al., 1990). Also, studies have explicitly tested the temporal frequency tuning of LGN and visual cortical neurons (Derrington and Lennie, 1984; Foster et al., 1985; Lee et al., 1989; Levitt et al., 1994) using drifting sinusoidal gratings—stimuli that rarely induce temporal modulation in the output of MT cells (Sclar et al., 1990; J. Anthony Movshon, personal communication). The stimulus used in the present study is different from those of most previous studies that have focussed on temporal properties of neurons because it is both random and dynamic. We will characterize the modulation of the neuronal response in both the time and frequency domains so that our results may be compared to those from other studies of temporal modulation.

4.1 Data Analysis Methods

All trials having a constant seed were grouped together and analyzed, but the values of the constant seeds were not stored, precluding stimulus reconstruction or the use of reverse correlation methods (McLean and Palmer, 1989; Bialek et al., 1991).

We analyzed 54 cells in three monkeys: 26 for E, 9 for J, and 19 for W. (For monkey E, 22 of the 25 cells recorded in the old experiments were used here, plus 4 cells recorded in later experiments on pairs of neurons.) Not all cells were recorded under all experimental conditions, so the number of cells involved in each analysis will be stated in the text. Cells with mean responses that changed by more than 100% during recording sessions were not included. In all computations, the PSTH for a set of spike trains was the average number of action potentials plotted as a function of time relative to stimulus onset and was computed at the same millisecond resolution as the original recordings—the data were not smoothed. The PSTHs shown in the figures,

however, were smoothed using an adaptive square window that was widened to include a criterion number of spikes. A one-sided estimate of the power spectral density of the PSTH was computed using the standard Fast Fourier Transform (FFT) algorithm and overlapping data segments with windowing (Press et al., 1988). To avoid biasing our statistics with the initial transient response, we restricted our analysis, except where noted, to the 1600 msec "sustained" portion of the spike trains that follows the 400 msec "transient" period beginning at stimulus onset, $t = 0$. Bursts of action potentials (consecutive spikes occurring with inter-spike intervals of 3 msec or less) can create excessive power at low frequencies for some neurons and therefore were replaced by single action potentials using the technique described in Chapter 7.

This analysis focuses on properties of sets of spike trains that were recorded from single neurons using the same dynamic dot sequence. Sets of responses for randomly seeded stimuli, with all other parameters held constant, are referred to as "control" data. Ideally, the trials in a set of control data will have the same statistical properties as the responses to any particular pattern, but will not show stimulus-locked modulation in the sustained period of the PSTH. An example of a set of control data is shown at the left in Fig. 2.1.

The following simple test for a violation of Poisson statistics was used to determine the presence of stimulus-locked modulation in a set of spike trains. The average firing rate was determined for the 1600 msec sustained period and taken to be the mean rate of a homogeneous Poisson point process. If the observed firing rate in any segment of the sustained period was improbably high at the 10^{-6} significance level, the response was considered to have stimulus-locked modulation. The stringent significance level reflects the inadequacy of the Poisson process to account for the refractory period, burst firing, and non-stationarities which are frequently found in spike trains. Of the 54 cells, 49 were found to respond to a dynamic dot sequence with significant temporal modulation based on 30 trials at $c = 0$. When run on 54 sets of control

data taken from the same number of cells in the same animals, The test for significant modulation yielded 2 false positives. Visual inspection of these false positives revealed that the spike rate changed slowly over the course of the trial and that this change was independent of the particular dot pattern.

4.2 Results of Temporal Analysis

A neuron presented with 90 different random dot stimuli at $c = 0$ produced an ensemble of responses (Fig. 2.1, left) that, except for the initial transient, can be approximated by a point process with a time-invariant mean rate, such as a homogeneous Poisson process modified by a refractory period (Bair et al., 1994). The right side of Fig. 2.1, showing 90 responses of the same cell to one *particular* $c = 0$ random-dot stimulus, reveals that the neuron's firing pattern was very tightly locked to the stimulus. Thus, much of the apparent randomness of the ensemble on the left is caused by the fact that a different random dot pattern was presented on each trial. The firing pattern on the right can be modeled to first order by a random process with a time varying mean rate, such as an inhomogeneous Poisson process. The time varying modulation, estimated by the post-stimulus time histogram (PSTH, bottom right), is characterized by narrow peaks, often produced by single action potentials occurring at precise instants across trials. A second neuron responded to a dynamic dot sequence at $c = 0$ with a much higher firing rate but was still highly modulated (Fig. 4.2).

After characterizing the temporal modulation in the time and frequency domains, we will contrast the patterned responses in Fig. 2.1 and Fig. 4.2 with the case of $c = 1$ stimulation, i.e. coherent motion, in which modulation is **not** present when an identical stimulus is repeated.

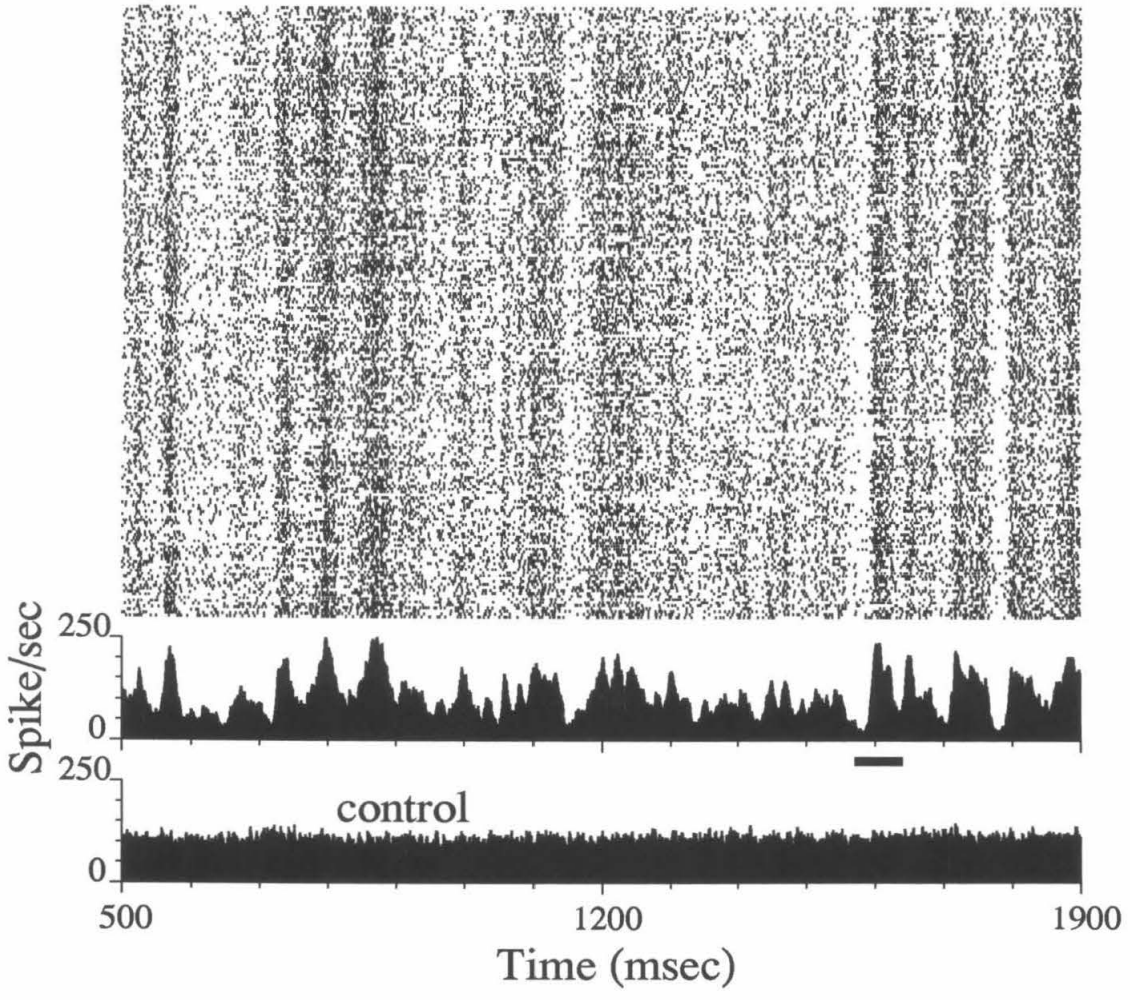

Figure 4.2: The neuronal response of cell e093 to $c = 0$ dynamic dot stimulation on 206 trials using a particular stimulus pattern. This neuron produces a highly modulated response, like the neuron in Fig. 2.1, but has a much higher firing rate (113 Hz, SD 15 Hz, 206 trials). While the responses of most cells contained occasional clearly isolated epochs of elevated firing rate, this cell never dips down to its background firing rate (2 Hz). The temporal precision of the response peak indicated by the thick bar below the PSTH is shown in Fig. 4.3, and the power spectra for both PSTHs are shown in Fig. 4.4. The lower, flat PSTH corresponds to 210 trials of control data. The adaptive window used to smooth the PSTHs was widened to include 40 spikes.

4.2.1 Precision and Reliability

For the 49 cells which had statistically significant stimulus-locked modulation (see Data Analysis Methods), we quantified the temporal precision of the spike trains using the standard deviation (SD) in time of the onset of periods of elevated firing such as those indicated by peaks in the PSTH at the bottom right of Fig. 2.1. This technique is similar to that used by Sestokas and Lehmkuhle (1986). The standard deviation measure will be referred to as temporal "jitter" and has a small value for a highly precise response. Occurrence times of action potentials were recorded with a 1 msec resolution. In the worst case, this discretization would add 0.5 msec to the temporal jitter. Other sources of error in the recorded time of action potentials relative to stimulus onset were small compared to the discretization error.

A peak in the PSTH corresponding to a period of significantly elevated firing probability (see Data Analysis Methods) was accepted as well enough isolated for analysis if an arbitrary point in time preceding the peak existed such that the mean time to the first spike in the response was greater than twice the standard deviation of the distribution of first spike times. For example, one statistically significant peak is marked by a thick line near 1740 msec in the PSTH at the bottom right of Fig. 2.1. Considering a period of 70 msec surrounding that peak, we measure the time from the beginning of the period to the first spike on each trial. The distribution of first spike times is shown in Fig. 4.3**A**. The SD, or jitter, is 3.3 msec and yet no action potential occurs for at least 25 msec prior to the response. The distribution achieved by this method is different from the shape of the peak in the PSTH which includes all the spikes, not just the first one in a response. In addition, by considering only the first spike in an isolated response period on each trial, our measurement is less likely to be biased by a refractory period or the inter-spike interval statistics for the neuron. For a few neurons, such as the one in Fig. 4.2, the significant peaks in the PSTH were not

well enough isolated to perform the jitter analysis for individual action potentials. In those cases, we searched for the first occurrence of a *pair* of spikes with less than a specified inter-spike interval (6 msec in Fig. 4.2). The distribution of the occurrence of these pairs is shown in Fig. 4.3**B**.

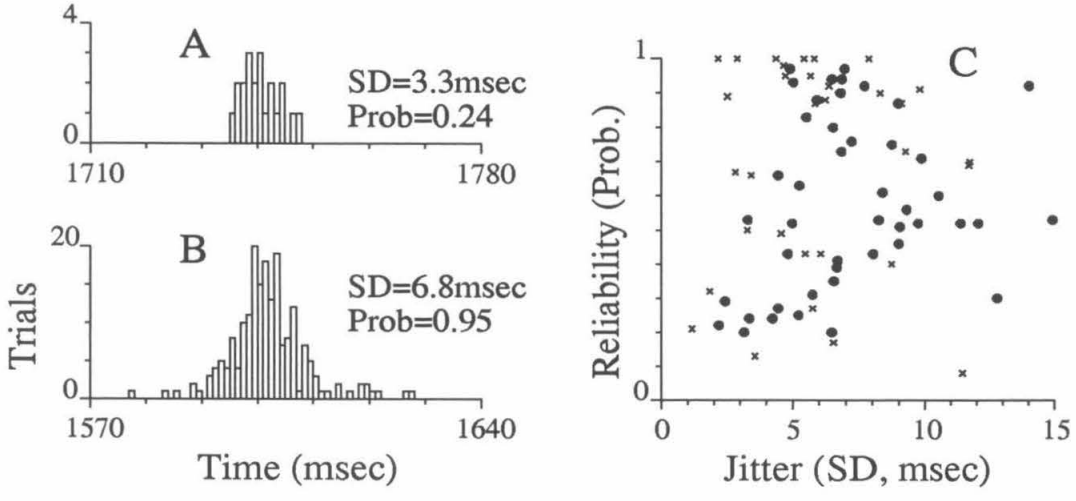

Figure 4.3: MT neurons are temporally precise on the order of milliseconds. We measured precision as the standard deviation (SD), or jitter, of the beginning or ending time of a period of elevated firing, e.g. the periods indicated by the thick lines below the PSTHs at 1750 msec in Fig. 2.1 and near 1600 msec in Fig. 4.2. **A.** The distribution across trials of the occurrence time for the first action potential in the response during the period 1710–1780 msec in Fig. 2.1. **B.** The distribution of the occurrence time of the first pair of action potentials fired within 6 msec in the time period 1570–1640 msec in Fig. 4.2. **C.** The jitter for the most precise response periods is plotted against the response reliability (probability) for 49 cells (solid circles). In 80% of cells, the minimum temporal jitter during the sustained portion of the response was less than 10 msec, with the smallest values near 2–3 msec. For comparison, crosses indicate the jitter of initial transients (present in only 32 of 49 cells). At least 30 repeated trials were used for each cell. Four points exceeded the horizontal scale, with jitter values of 16, 19, 27 and 69 msec.

Reliability was measured as the probability that a response occurred during the periods described above. For the response period analyzed in Fig. 4.3**A**, the cell

responded (fired any number of action potentials within the response peak) on 24% of trials (see Fig. 2.1, right, for spike trains). The same cell had other, more reliable responses—the reliability for the response near 1000 msec was 84%. Within the response period marked on the PSTH for the neuron in Fig. 4.2, two spikes with inter-spike interval ≤ 6 msec occurred on 95% of the trials. The distributions shown in Fig. 4.3**A** and **B** correspond to the most precise responses for the spike trains shown in Fig. 2.1 (right) and Fig. 4.2. For all 49 cells, the scatter plot in Fig. 4.3**C** shows reliability versus jitter for the most precise response during the sustained period (filled circles) and for the initial transient, present in only 32 of 49 cells (crosses). For 80% of cells, the most precise response during the sustained period had jitter less than 10 msec and in some cases the jitter was as small as 2–4 msec. The initial transients typically had less jitter than the most precise sustained period response.

4.2.2 Frequency Profile

The previous analysis focussed only on the most precise period of the responses in the time domain, but we now examine the entire sustained period response in the frequency domain. Temporal frequency profiles of the responses of the MT cells were computed as the power spectra of the PSTHs for $c = 0$ stimuli. Spectra are shown in Fig. 4.4 for the PSTHs in Fig. 2.1 (right) and Fig. 4.2. The power spectra are consistent with the notion that the cells act as low pass filters for the white noise dot stimulus; however, it is important that these profiles are not mistaken for temporal frequency tuning curves (see Discussion). There were no systematic peaks in the spectra at particular frequencies; for example, there was no stimulus refresh artifact as the dots were plotted asynchronously. Note that in Chapter 6 stimulus artifacts were observed for $c = 1$ stimulation.

To compute an upper cutoff frequency for individual cells, we compared the power

spectrum for a particular stimulus pattern to that for a PSTH computed from control data in which all trials resulted from different $c = 0$ stimuli. The cutoff frequency was taken to be the lowest frequency at which the control power spectrum intersected the response power spectrum. The histogram of cutoff frequencies for 22 cells from monkey E reveals a range of values from 0 to 150 Hz (mean $\pm$ standard deviation $= 58 \pm 38$ Hz, Fig. 4.4 bottom). Response data and control data from the *same* cell was only available for monkey E, but data recorded for individual *coherent* ($c = 1$) motion stimuli for J and W served as a control because, as reported in the next section, modulation was virtually absent for coherent motion. The distribution of cutoff frequencies for nine cells from monkey J had a mean (46$\pm$10 Hz) that was not significantly different than that for E (t-test, $p = 0.19$), while the mean (23$\pm$12 Hz) for six cells from W was significantly lower than that for both E and J (t-test, $p < 0.005$).

The distribution of cutoff frequencies shown in Fig. 4.4 is consistent with our analysis of data from other experiments, using a different display, in which 73% of cells (22 of 30) in monkey J showed a peak in their power spectrum at 40 Hz when a coherent dot stimulus ($c = 1$) was presented in frames at 40 Hz. In an experiment using moving bars, 24% of cells (12 of 49) in a fourth monkey (R) showed peaks in their spectra at 60 Hz when the bar moved on a 60 Hz frame-refresh monitor.

Because the autocorrelation of a function is the Fourier transform of its power spectrum, these results may be interpreted in the time domain from the autocorrelation of the PSTH. The autocorrelations (computed after subtracting the means from the PSTHs) displayed a single peak before falling to zero with a width at half-height of 36 $\pm$20 msec (means for individual animals were: E 25$\pm$9.2, J 33$\pm$16, W 50$\pm$24). Note that both the autocorrelation function and the power spectrum are computed from the PSTH, and this is different from computing these functions for individual trials and averaging afterwards (as done by Bair et al., 1994).

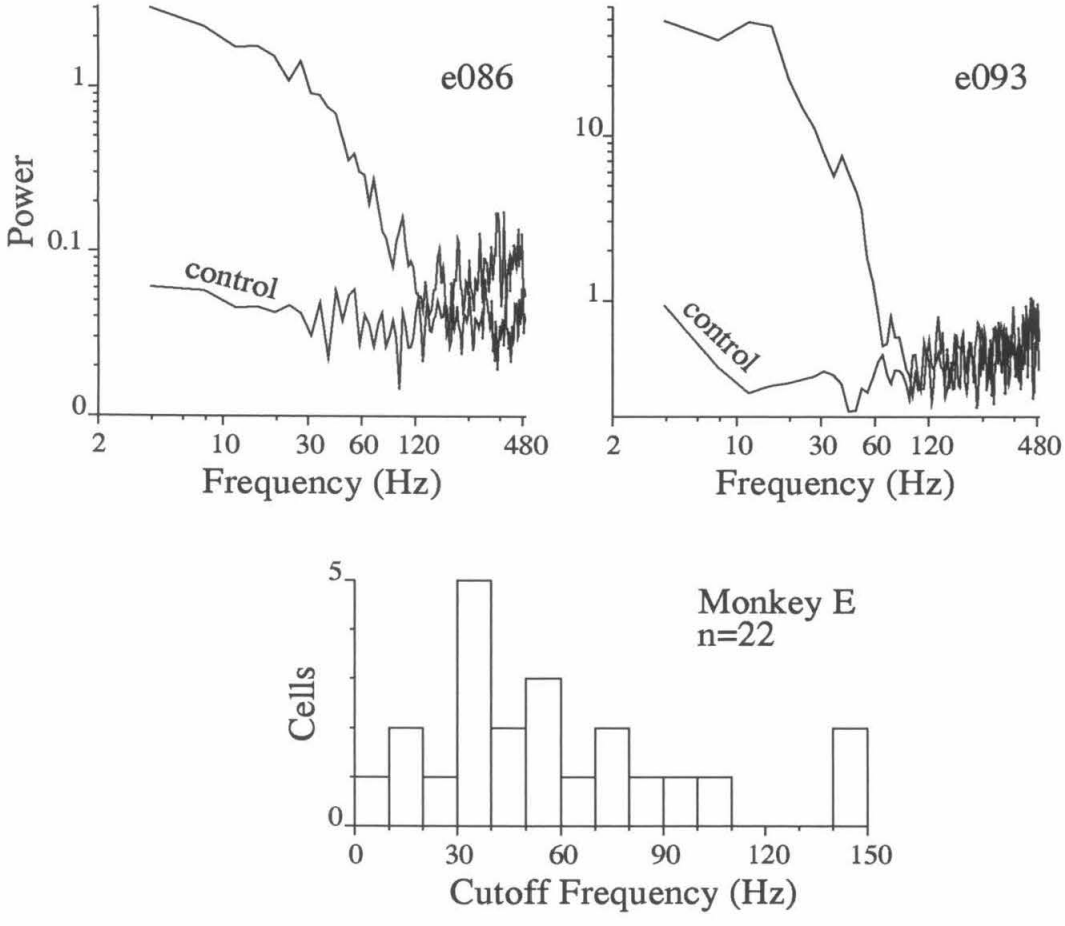

Figure 4.4: Temporal frequency cutoffs for area MT cells in response to dynamic random dots. The upper panels show the power spectra of the PSTHs for the neuronal responses shown in Fig. 2.1 (right) and Fig. 4.2. We defined the cutoff frequency for each cell as the lowest frequency at which the control spectrum intersects the response spectrum. The lower panel shows the distribution of frequency cutoffs for 22 cells from monkey E. Temporal frequency cutoff (defined as inducing 50% of the optimal response) distributions for MT cells based on sine wave gratings are concentrated below 32 Hz (Michael J. Hawken, personal communication). A cell-by-cell comparison of the temporal frequency profile based on sine waves and the power spectra in response to the random dot stimulus would be revealing but has never been made. (Monkey E was the only animal for which control data was available, see text.)

4.2.3 Response to Coherent Motion

The presence of temporal modulation depended on the motion coherence of the stimulus. While it was apparent for low coherence stimuli as shown in Figs. 2.1 and 4.2, it was absent for highly coherent motion, *i.e.* $c = 1$ (Fig. 4.5). We defined a measure M of the overall modulation strength based on the power spectrum of the PSTH. M was the integral of the power spectrum in the 4–30 Hz band divided by the mean spike rate across the PSTH. As shown in Fig. 4.6 and by the range of cutoff frequencies in Fig. 4.4, temporal modulation showed up as excessive contributions to this frequency band. In monkeys J and W, repeated stimulation using a particular dot pattern was performed at higher coherence levels, $c = 0.5$ and 1.0, in addition to $c = 0.0$. In both animals, the modulation strength M was not significantly different at $c = 0$ versus $c = 0.5$ ($p > 0.20$, paired t-test). Yet, for both animals, M was significantly less at $c = 1$ compared to $c = 0.5$ (statistical significance: $p < 0.005$ for monkey J, $p < 0.05$ for monkey W).[1]

We note that M was not significantly correlated with spike rate ($r = 0.09$, $p = 0.50$, M vs. log of spike rate) nor with the diameter of the stimulus aperture, which was optimized for the receptive field of each cell ($r = 0.15$, $p = 0.31$).

4.2.4 Discussion

We have observed that a dynamic dot stimulus can produce periods of precisely timed, stimulus-locked modulation in the spike trains of neurons from area MT. While it is common practice to seek the stimulus that causes the "largest activity" (Lettvin et al., 1959), we have sought those periods in which the stimulus caused the most precise activity to form an estimate of the temporal precision with which the cortical network

[1] Because $c = 1.0$ stimulation was not used for all cells, these computations were based on nine cells in monkey J and six cells in monkey W. In all computations, the average M computed for control data from *other* cells within the same animal was subtracted.

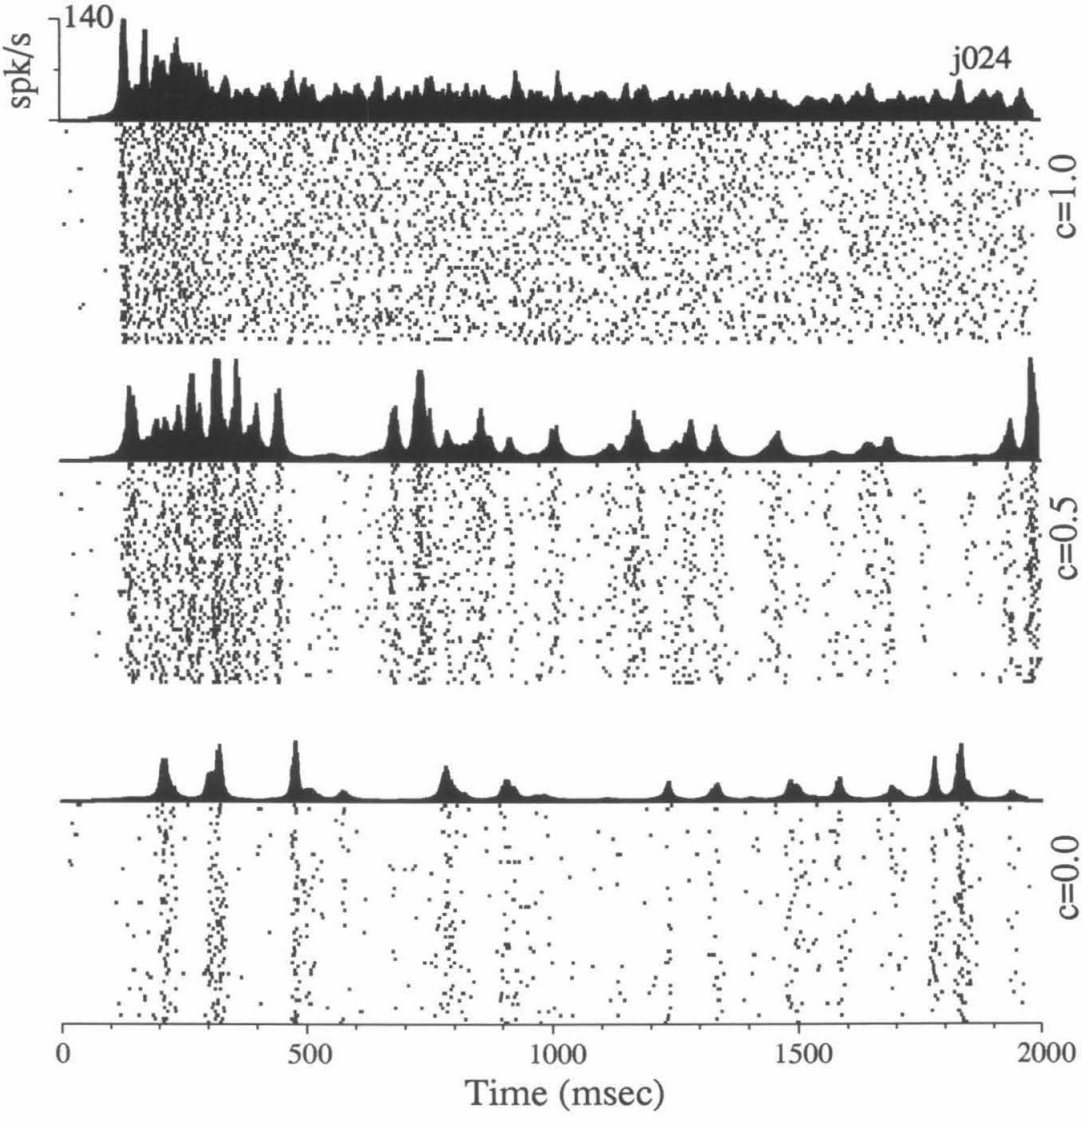

Figure 4.5: Temporal modulation disappears for highly coherent stimuli. The spike trains and PSTHs demonstrate that the stimulus-locked temporal modulation present for incoherent motion ($c = 0$) and for partially coherent motion ($c = 0.5$) was virtually absent during the sustained period of the response to coherent motion ($c = 1$). This suggests that temporal dynamics beyond those found in rigid translation are necessary to induce a specific and unique time course in the spike discharge pattern.

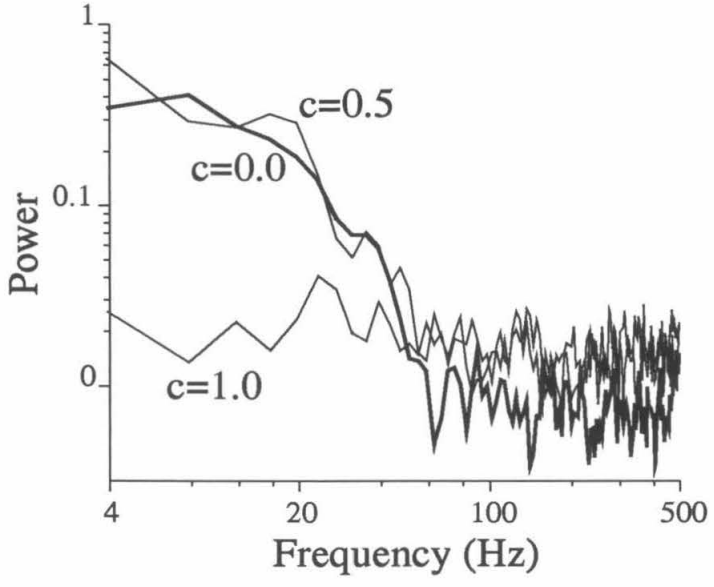

Figure 4.6: Power spectra of the PSTHs in Fig. 4.5. The temporal modulation, reflected by excess power at frequencies below 60 Hz, was similar for $c = 0$ and $c = 0.5$ but was absent for coherent motion. The PSTHs were expressed as instantaneous firing probabilities and the resulting power spectra were normalized by the mean firing rate. Our measure M of temporal modulation was the integral of the scaled spectrum in the 4-30 Hz band. Fifty-nine spike trains were included at each c.

can trigger an action potential within a neuron. Using the standard deviation measure of temporal jitter, we found that 80% of cells were capable of responding with jitter less than 10 msec, and the most precise responses had jitter less than 2 msec. The reliability of the cell, its probability of contributing to a peak in the PSTH on a particular trial, varied widely from 0 to 1 (Fig. 4.3).

The output frequency profiles of the averaged responses of MT cells (Fig. 4.4) are all low-pass with a broad range of cutoffs, some above 100 Hz, while the visual input provided by the incoherent dynamic dot stimulus has a flat temporal frequency spectrum.[2] Studies that used drifting sinusoidal gratings to determine the temporal frequency sensitivities of LGN and visual cortical neurons (Derrington and Lennie, 1984; Foster et al., 1985; Lee et al., 1989; Levitt et al., 1994) found cutoff frequencies distributed lower than the distribution of cutoff frequencies reported here; however, we emphasize that the frequency profiles of the responses reported here should **not** be interpreted as temporal frequency tuning curves. We do not know, for example, whether power near 50 Hz in the response is caused by 50 Hz components in the stimulus or arises from a computation such as squaring or rectifying a 25 Hz input component.[3] This could be tested if the seeds to the stimuli were known. However, it is not unreasonable to assume that area MT processes information that has frequency components in the 100's of Hz since temporal order, i.e., apparent motion, can be perceived for temporal separations down to 3 msec for a pair of optimally positioned bars (Westheimer and McKee, 1977). In addition, apparent vernier offsets can be perceived for temporal intervals as short as 1.9 and 2.2 msec (Burr, 1979). Even under the assumption that all inputs to area MT have relatively low temporal cutoff frequencies, high frequency signals may still be reconstructed from the spatial distri-

[2]In practice, the approximate 150μsec lifetime of the dots is so short that the departure from a flat temporal spectrum is not significant to the monkey visual system.

[3]Squaring a sine wave causes frequency doubling, and rectification introduces even higher harmonics.

butions of the inputs, in analogy to spatial hyperacuity. But there is evidence that precise and high frequency responses are carried by single neurons earlier on. Fast initial transients (having 50–100 Hz oscillations) are observed in magnocellularly derived responses in V1 (Maunsell and Gibson, 1992) and, in the cat, retinal ganglion cells can lock their output to 100 Hz flicker (Eysel and Burandt, 1984).

The stimulus-locked temporal modulation that accompanies incoherent and partially coherent motion is not present for coherent motion, even though the same dots are flashed at the same time and location from one presentation of the coherent stimulus to the next and in a manner such that only a single dot is likely to be within a V1 sub-unit of an MT receptive field at any time (within 50 msec or longer). In terms of the output, spatial inhomogeneities in the MT receptive field (so called "hot spots"), if they exist, are not apparent for this type of rigid pattern translation. Snowden et al. (1992), using a smaller and more dense stimulus (3° aperture), estimated that 90% of MT cells did not modulate to a moving random dot stimulus; however, because the motion coherence of their stimulus is approximately equivalent to that of a $c = 0.96$ stimulus here, their results are consistent with the lack of modulation that we find at $c = 1.0$. For coherent motion in the neuron's preferred direction, we return to the simplest statistical description of the output of the cell: the mean spike count during the stimulus. In this case, the timing of individual action potentials may be governed by noise, but if more careful experimentation reveals that the spikes are precisely locked to internal fluctuations in the cortical network, we may realize, as Barlow (1972) observed for individual nerve cells, that "...their apparently erratic behavior was caused by our ignorance, not the neuron's incompetence." In Chapter 5, we study the response to coherent motion more carefully and find evidence for modulation that is not directly stimulus locked, and correlation between the signals of nearby neurons.

What we find remarkable about these data is not primarily their relationship to

motion and motion energy models (Adelson and Bergen, 1985; Britten et al., 1993) but what they might imply about the significance of individual action potentials in cortical cells. It is a challenge to theories of cortical processing (Stevens, 1994) to explain how the observed low probabilities of synaptic transmission in cortical brain slices (Smetters and Nelson, 1993; Thomson et al., 1993; Allen and Stevens, 1994), combined with single synaptic contacts among pyramidal cells in mammalian cortex (Freund et al., 1985; Andersen, 1990; Gulyas et al., 1993), can give rise to these highly reproducible spike patterns in cells 7 to 8 synapses remote from the sensory periphery over a two hour long experiment in a behaving animal (see also Lestienne and Strehler, 1987; Abeles et al., 1993). Mainen and Sejnowski (1995) report that sustained current injection into cells in neocortical slice leads to a variable spike response, while repeated stimulation with a particular white noise current evokes a highly reliable spike pattern. Their findings in cortical cells suggest that the spike triggering mechanism itself is capable of accurately encoding temporally modulated input into spike trains, possibly providing the biophysical substrate of our results. However, it remains to be determined to what extent the rapid temporal modulation reported here carries any detailed information of behavioral significance.

In Chapter 6, we show that a motion energy model (Adelson and Bergen, 1985), when driven by the dynamic dot stimulus, can produce precise and high frequency temporal modulation. However, the model does **not** attempt to represent sources of jitter that might affect the neuronal computation and, therefore, does not explain how the transmission from retina to area MT remains so crisp. The argument presented in Chapter 6 will depend upon the connection, derived in the final section of this chapter, between the temporal frequency content of the signal and our precision measure.

4.3 Temporal Precision in the Frequency Domain

This section provides an analytical expression that links the measure of precision in the time domain to the frequency domain analysis. This analysis considers the special case in which the probability of firing an action potential is described by a half-wave rectified sine wave of a particular amplitude. The analysis predicts that even if only low frequencies are passed by the system, the appropriate normalization, consisting of truncation and rescaling, can lead to highly precise responses.

The probability density for the first spike time that is generated from an inhomogeneous Poisson process with mean rate $\lambda(t)$ is

$$f(t) = \frac{\lambda(t) e^{-\int_0^t \lambda(\tau)\, d\tau}}{1 - e^{-\int_0^\infty \lambda(\tau)\, d\tau}} \qquad (0 \leq t). \tag{4.1}$$

The numerator is related to the product of the probability that a spike occurs at time t and the probability that no spike occurs before time t. The denominator is related to the probability that at least one spike occurs at any time, i.e., one minus the probability that a spike never occurs. If $\lambda(t) = \lambda_o$, $f(t)$ is simply the decaying exponential distribution,

$$f(t) = \lambda_o e^{-\lambda_o t}, \tag{4.2}$$

which describes the inter-spike interval distribution of a homogeneous Poisson process. We consider the case where the mean rate, $\lambda(t)$, of the inhomogeneous process is given by one half period of a half-wave rectified sine wave having frequency f_o and producing on average N spikes per lobe, i.e,

$$\lambda(t) = N\pi f_o \sin(2\pi f_o t) \qquad (0 \leq t \leq 1/(2f_o)). \tag{4.3}$$

Substituting this into Eqn. 4.1 yields the probability of first spike time

$$f(t) = \frac{N\pi f_o \sin(2\pi f_o t)}{1 - e^{-N}} e^{-\frac{N - N\cos(2\pi f_o t)}{2}}. \tag{4.4}$$

This probability distribution is plotted in Fig. 4.7 for the case $N = 5$ spikes and $f_o = 32$ Hz. This distribution is designed to model those in Fig. 4.3 A and B. As in Fig. 4.3, the temporal precision is computed here as the standard deviation of the distribution of first spike times. The standard deviation depends upon the frequency and amplitude, i.e., firing rates, of the sine wave in Eqn. 4.3. Fig. 4.8 shows the precision that can be achieved for varying frequencies (f_o) and varying spike rates (N). In Chapter 6, we show that given the frequencies passed through a quadrature energy model, and depending on the type of normalization, one can achieve highly precise spike trains.

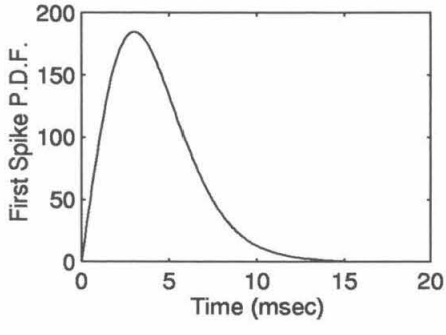

Figure 4.7: The probability density function for the first spike time for one lobe of a half-wave rectified sine wave of frequency $f_o = 32$ Hz and spike density $N = 5$. The standard deviation of the probability density function is 2.3 msec, corresponding to the most precise events observed in the stimulus locked temporal modulation of the MT spike trains.

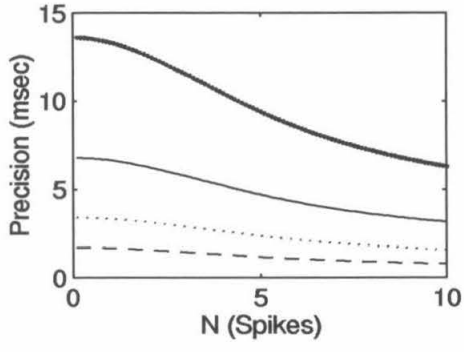

Figure 4.8: The temporal precision (standard deviation) of distributions like that shown in Fig. 4.7 (Eqn. 4.4) are plotted for $f_o = 8$, 16, 32, and 64 Hz (from top to bottom) as a function of mean spikes per event (N).

Chapter 5

The Response to Coherent Motion

Coherent motion, i.e., $c = \pm 1$, is usually the strongest stimulus for the MT neurons in terms of elevating or suppressing mean spike rate during the sustained period of the response, and the neuronal discharge evoked by this stimulus condition is the most likely to resemble that evoked by a natural environment. Results of the previous chapter suggest that the coherent motion stimulus caused little, if any, stimulus-locked modulation during the sustained period of the response, in marked contrast to the incoherent motion stimulus. Following this observation, the $c = \pm 1$ spike trains were analyzed as a separate group, with the idea that any remaining temporal modulation could be intrinsic to the cortical processing and not due to stimulus variations.

This chapter begins by describing the occurrence of excessively long inter-spike intervals, or *gaps*, that occur in the brisk response to preferred motion and excessively dense clusters, or *bursts*, of spikes that occur in the suppressed response to null motion. (The bursts described here occur at a longer time scale than those described later in Chapter 7.) In the second section of this chapter, a novel cross-correlation technique reveals that the gaps tend to occur at similar times in pairs of neurons that were recorded from a single electrode in a previous study (Zohary et al., 1994). Common

inhibition is discussed as a possible mechanism underlying the gaps.

In the absence of stimulus induced modulation, the $c = 1$[1] data are also ideal for estimating the variability of the spike trains. A renewed interest in spike train variability (Softky and Koch 1993; Shadlen and Newsome, 1994; Usher et al., 1994; Bell et al., 1995) has in part been sparked by a different analysis of this MT database (Softky and Koch, 1993). The final section of this chapter compares the variability of the inter-spike intervals for $c = 1$ stimulation to that for incoherent stimulation. Variability of spike count is also examined, and a statistical model reveals that the presence of long gaps in the spike train can create a power-law dependence of the variance upon the mean spike count, consistent with previous reports from cat and monkey visual cortex. The chapter ends with a discussion of variability and speculation about the nature of the signal that is supplied to the axon initial segment under coherent motion stimulation.

5.1 The Symmetry of Gaps and Bursts

The response of an MT neuron to a random dot stimulus moving coherently in its preferred direction can roughly be broken into a transient and a sustained period. The transient period can be highly variable, but is marked by adaptation which leads to a sustained firing rate within the first 600 milliseconds (Allman et al., 1985). Early in the sustained period, typically 600–900 msec post-stimulus onset, we have observed excessively long ISIs, or *gaps*. Fig. 5.1 (top) shows an example of a neuron's brisk response to preferred motion being interrupted by gaps that are roughly 100 msec long and occur predominantly in the second quarter of the 2 sec response. This is not accounted for by the slow, steady adaptation (presumably due to potassium currents)

[1] At $c = -1$, the firing rate is often too low for analysis.

which is observed under current injection in neocortical pyramidal neurons, e.g., the RS_1 and RS_2 neurons of Agmon and Connors (1992).

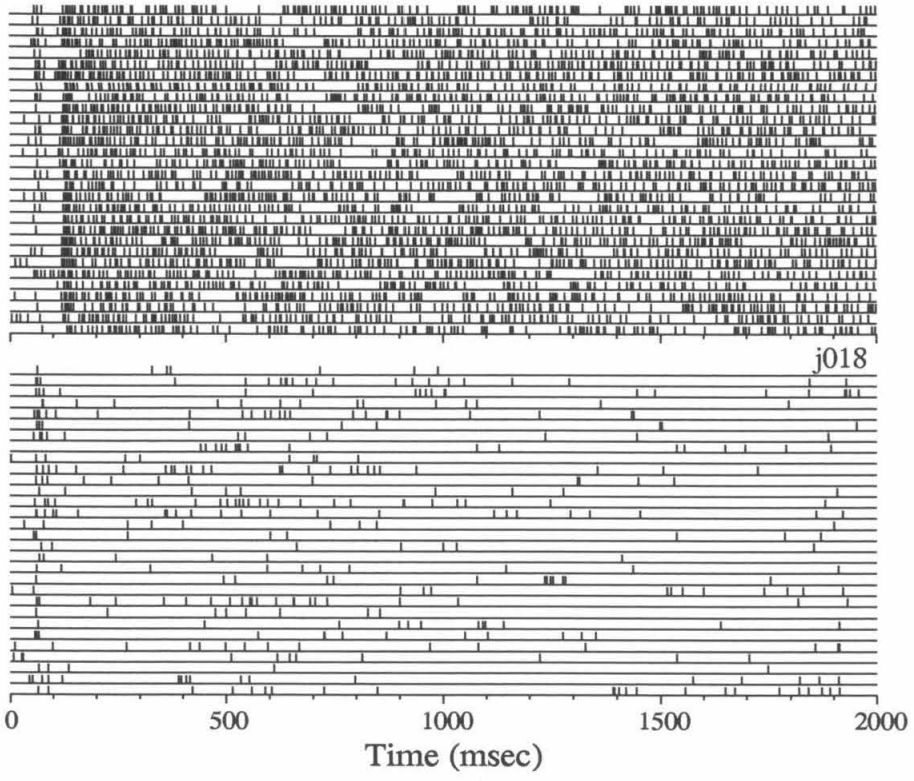

Figure 5.1: The appearance of long inter-spike intervals near 600–900 msec in the brisk response to preferred direction motion (top) is mirrored by clusters of action potentials in the weak response to null direction motion (bottom).

The bottom of Fig. 5.1 shows a symmetrical phenomenon which occurs for null direction motion. Clusters of action potentials occur during the null stimulus at nearly the same time as the gaps did for preferred motion. In some neurons, the clusters are even more highly characterized and restricted in time than the gaps, but over all, the gaps appear to be a more robust feature. Two additional examples, Fig. 5.2 and Fig. 5.3, show the diversity of the gaps and bursts.

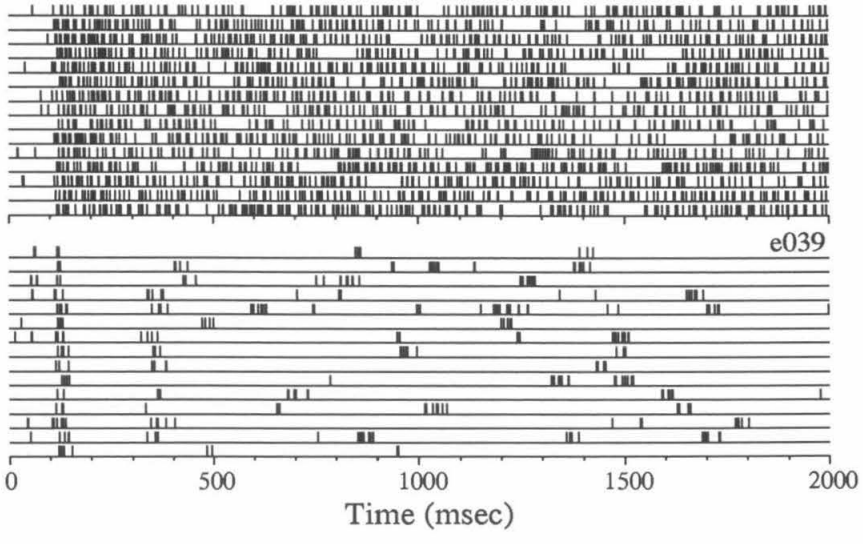

Figure 5.2: The bursts in response to null, $c = -1$, motion (bottom) can have a duration that is as characteristic as that for the gaps in the response to preferred motion (top), as shown here for cell e039.

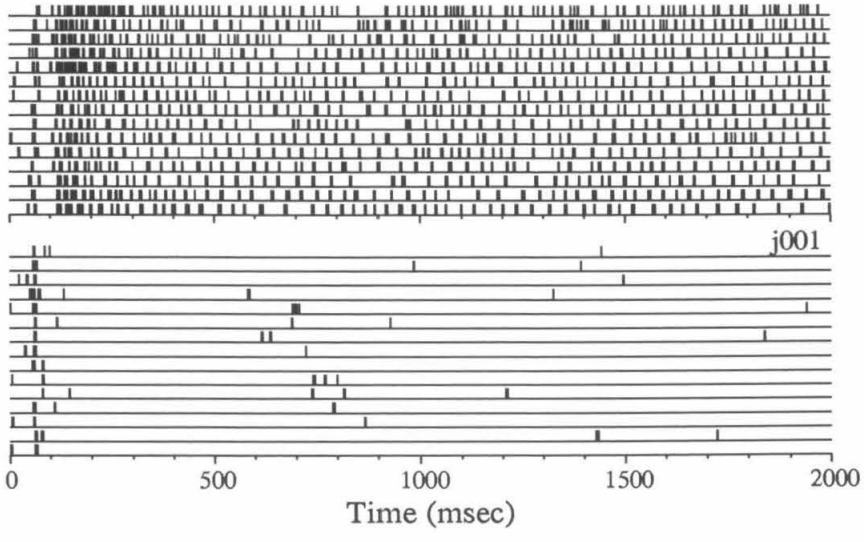

Figure 5.3: The bursts (bottom) and gaps (top) occur at roughly the same time, 600–900 msec for cell j001, however, there is no cell-by-cell correlation across the database. The bursts rarely are more localized in time than shown here.

We analyzed the 84 neuron subpopulation of the 216 neuron database for which both $c = 1$ and $c = -1$ stimuli were presented and for which the spike rate in response to $c = 1$ was at least 20 Hz (since gaps would not be detectable at low firing rates). For each neuron, we determined the peak time of occurrence of gaps during preferred trials and the peak time of bursts during null trials. Because the long ISIs can be hidden in a post-stimulus time histogram (PSTH) by the surrounding high firing rate (particularly if adaptation has not subsided), we converted the spike trains to their *interval* representation. The interval representation at each millisecond is simply the length of the ISI encompassing that sampling point. This representation (an example of which is shown later in Fig. 5.9) looks like the a city skyline where skyscrapers are long ISIs. IPSTHs (interval PSTHs) were computed for preferred responses by summing the interval trains, and regular PSTHs were computed for the null responses by summing the spike trains. Both are shown smoothed with a Gaussian of $\sigma = 40$ msec in Fig. 5.4 for the spike trains in Fig. 5.1.

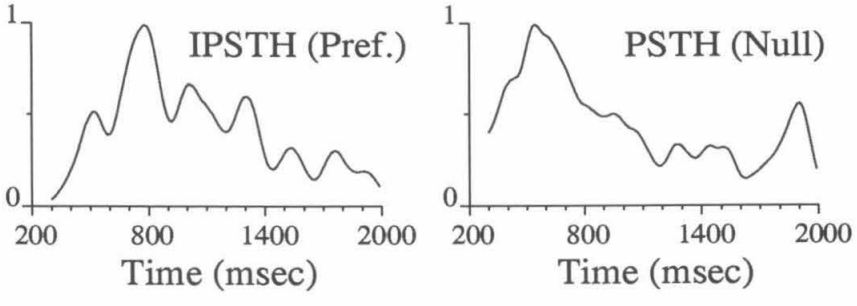

Figure 5.4: The normalized interval PSTH (IPSTH) for preferred motion and the regular PSTH for null motion. The peak of each plot is summarized in Fig. 5.5. for the subpopulation of neurons studied here.

Frequency histograms of the time of the peak in the preferred IPSTH and the null PSTH are shown in the left and right, respectively, of Fig. 5.5. The left panel of Fig. 5.5 shows that 76% of cells have their longest ISIs before 1200 msec, with a

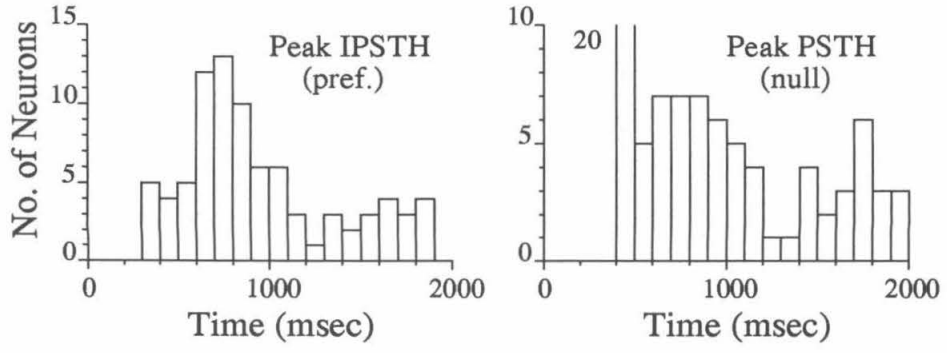

Figure 5.5: Across the population of neurons, the longest ISIs occur more often between 600–900 msec for preferred direction motion (left). This is mirrored by the densest occurrence of action potentials at nearly the same time for null motion (right). The dip at 1200 msec appears in both plots also. In the null motion case, twenty neurons had highest firing rates closer to the beginning of the trial (open bar, right). In some cells, the distinction between the initial transient and the burst response becomes ambiguous.

preference for 600–900 msec. The right panel indicates that the clusters of spikes in the null responses also occur most densely before 1200 msec, with a preference for 600–900 msec. Analysis of the null data is complicated by a trend for the firing rate to decay throughout the first quarter of the trial. Thus, 13 of the 20 cells in the first bin actually had their highest firing rate earlier than 400 msec, but this portion of the response is dominated by initial transients. The similarity of the time courses of the gaps and the bursts across the population lead us to describe these phenomena as symmetrical; however, we find that there is no correlation ($r = 0.06$, $p = 0.66$) on a cell-by-cell basis between the times of the peaks in the IPSTH and PSTH.

Gaps occurred in the previous examples for $c = 1$ stimulation, a condition which does not impose stimulus-locked temporal modulation on the spike train. Fig. 5.6 shows data from other experiments[2] in which the dot stimuli were shown in syn-

[2]These data were recorded in the "split-field" experiments (unpublished data of K. H. Britten and W. T. Newsome).

chronous "movie" frames at 40 Hz. This induced a structured firing pattern composed of bursts of spikes with 25 msec center-to-center spacing. The gaps tend to destroy two cycles of response, but very rarely three—indicating that the mechanism causing the gaps lasted less than 75 msec for this cell. The ISI distribution in Fig. 5.7 for these data shows more directly that the gap duration tends to cluster below 75 msec (3 cycles $*$ 25 msec/cycle, minus about 8 msec to account for the lengths of the bursts) with some around 50 msec, and only very few near 100 msec. This suggests that, if created in a controlled manner, precise temporal modulation may provide a tool for revealing properties of neurons. For example, it would be interesting to induce bursts every 80 msec to test whether the gaps would remain visible by destroying the bursts or whether they would disappear or possibly lie between the bursts.

Our final example violates two rules of conduct: (1) do not argue based on a single neuron, and (2) do not analyze data from an injured neuron. Nevertheless, Fig. 5.8 shows that in the early trials (within each block, trials are shown in the order collected, with the first trials at the top) long gaps appear by 500 msec into the trial. On later trials these gaps are not present. Such a drastic non-stationarity is often taken as an indication that the cell, or its inputs, were damaged by the electrode; however, we cannot definitely rule out that these changes might be related to normal cortical function. The trials in the lower block were recorded at $c = -1$ and, in the original experiment, were interleaved with the trials of the upper block, so it is clear that the cell maintained its direction selectivity. For example, if only the data from 0–300 msec are considered, the cell is strongly and consistently directional. This example is consistent with the notion that the gaps arise from a mechanism that is independent of the one that carries the signals that account for the directional selectivity of the neuron. The signals providing direction selectivity are likely to arise somewhat directly from V1 inputs, while the signal creating the gaps may be generated by connections intrinsic to area MT. Of course, this is speculation, and

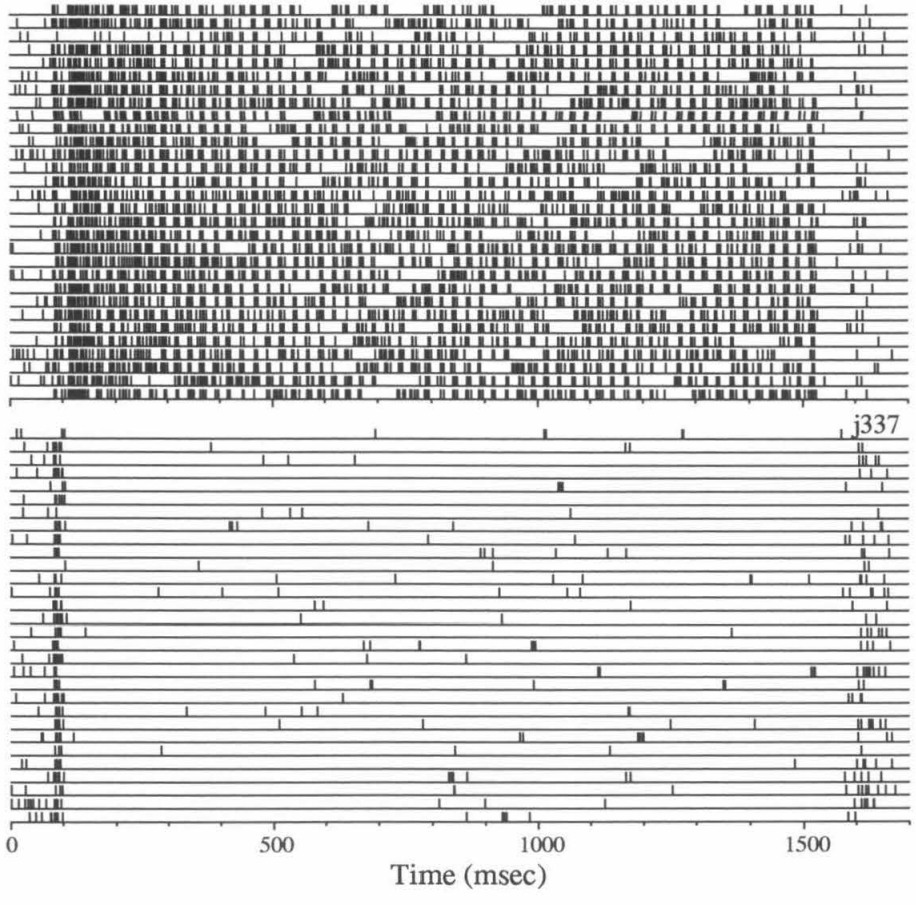

Figure 5.6: In the "split-field" experiments, 1.5 sec duration coherent dot stimuli were presented in frames at 40 Hz. This induced 40 Hz modulation in 22 of 30 cells studied. Bursts and gaps appeared under these experimental conditions as well, even though the temporal structure was under the control of the stimulus, unlike the previous $c = 1$ examples. Superposed on a known firing pattern, the time course of the gaps may be better characterized. The ISI distribution for the response to preferred motion (top) is shown in Fig. 5.7. (Data for ecode 5500—full field—was merged from left, right, and transparent files for cell j337.)

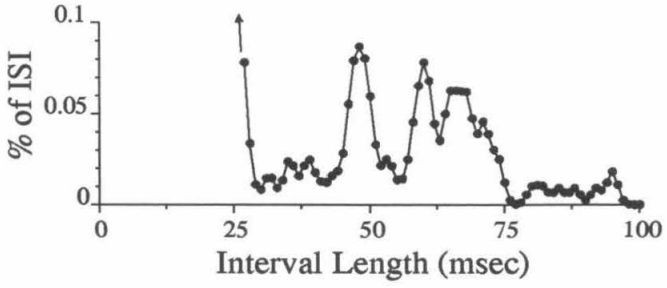

Figure 5.7: The ISI histogram (smoothed with a Gaussian of $\sigma = 1$ msec) shows the typical length of the gaps for cell j337 (Fig. 5.6, top). ISIs taken from t=300–1500 msec.

it is possible that the observed non-stationarity has nothing to do with the gaps in healthy neurons.

It would be particularly interesting to know if the gaps, and bursts, are correlated across neurons. If so, this could lead to identifying a mechanism to account for the correlation that is known to exist in spike counts from pairs of MT neurons (Zohary et al., 1994). In the next section, we look at the correlation between pairs of neurons with an analysis that emphasizes the gaps.

5.2 Interval Cross-correlation

If inhibition causes the gaps in the response to preferred motion, based on what is known about the architecture of inhibitory interneurons in cortex (Berman et al., 1992), it is likely that the gaps are correlated between nearby MT neurons. For instance, it was suggested early on (Somogyi et al., 1982) that axoaxonic interneurons, known to project to the axon initial segment of hundreds of pyramidal neurons within a 100–200μm diameter cylinder, would be useful for synchronizing the output of pyramidal neurons, and this has been demonstrated in models (Lytton and Sejnowski, 1991).

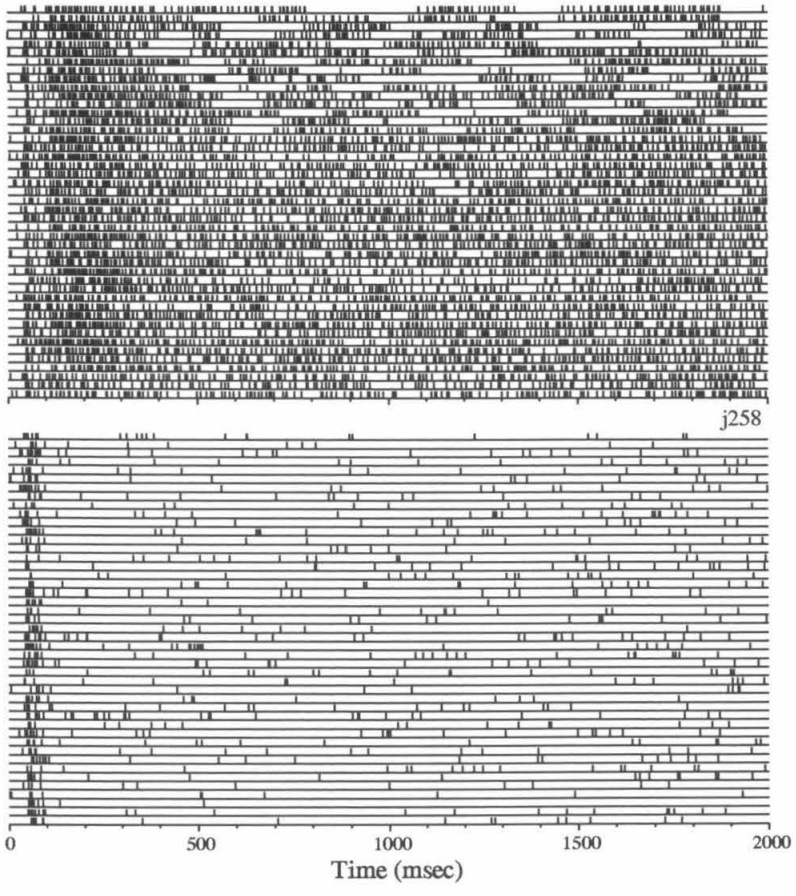

Figure 5.8: The discharge from an apparently injured neuron (j258) displays gaps that occur near the usual time in the first few trials recorded but occur less frequently on later trials. (The trials are displayed in the order recorded, but preferred (top) and null (bottom) blocks were interleaved.) If only the first 300 msec of the response is analyzed, the cell is strongly and consistently directional, in spite of the waning of the gaps.

To test this hypothesis, we analyzed spike trains from pairs of MT neurons that were recorded in a previous study (Zohary et al., 1994). These cells were recorded simultaneously using a single extracellular electrode. The standard motion discrimination paradigm was used, but the stimulus was generated on a 60 Hz raster scan device.

Standard cross-correlation analysis of extracellular spike trains is equivalent to representing the action potentials as 1's in an array of 0's, typically sampled at 1 msec or finer resolution, and computing the discrete cross-correlation function. This method is ideal for detecting relationships due to common sources of excitatory input, but common inhibition is notoriously more difficult to detect (Perkel et al., 1967; Aertsen et al., 1989). This is because at a millisecond resolution with typical firing rates of 20 spikes/sec in two neurons, it is unlikely that spike trains are cross-correlated at random; the probability of a spike in neuron 2 lining up with a given spike in neuron 1 is 1/50. Such low chance probabilities makes even a weak correlation in the timing of action potentials stand out above the background. However, the probability of there being no action potential from neuron 2 at a time when neuron 1 did not fire is 49/50. Thus inhibition can be very difficult to detect (and almost impossible to detect if it acts for only a short time).

We have developed a technique that can isolate and emphasize the contribution to the cross-correlation made by long ISIs. To emphasize the presence of long ISIs, the spike trains are transformed into interval trains as described above and shown in Fig. 5.9. This removes some of the information regarding the precise occurrence times of action potentials. For example, if two spike trains are perfectly regular at 40 Hz, the interval trains will be flat. The interval cross-correlation (ICC) between each pair of interval trains is computed and averaged over all trials, and the average shift predictor is subtracted. Fig. 5.10 shows ICCs (thick lines) for two different pairs of neurons. In 17 of 31 pairs (55%), there were peaks in the raw ICC that were at

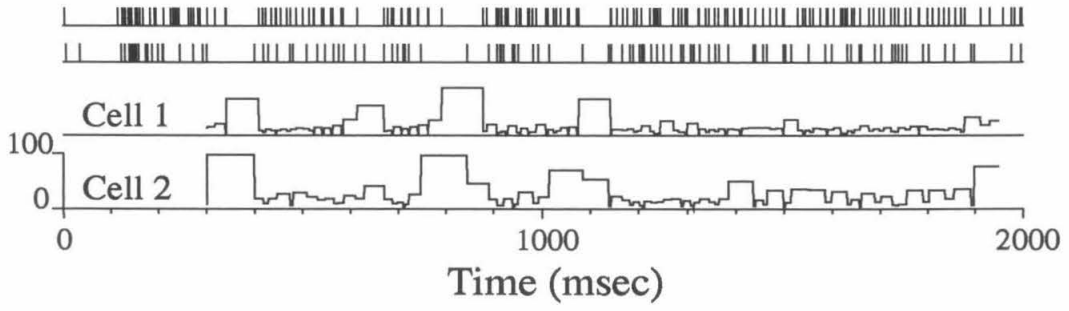

Figure 5.9: Simultaneously recorded spike trains (upper) and their interval representations (lower, amplitude in msec). Long ISIs transform to large boxes. The longest intervals seem to occur at similar times. This correlation is evident in Fig. 5.10 (Recording site emu084).

least 4 standard errors above the level of the shift predictor. The peaks were on average centered (mean 4.3 msec, SD 54 msec) and had mean width-at-half-height of 139 msec (SD 59 msec). The negative-going lobe in Fig. 5.10 (left) is caused by dense firing following the long ISIs in one of the cells. This occurs in roughly 6 cases, and is consistent with the notion that one cell fires above its mean rate directly following the gap, as would be the case for post-inhibitory rebound.

To isolate the cause of the peaks, the *long* intervals in the trains were set to the mean of the *short* intervals. The long intervals were defined as those that accounted for 30% of the duration of the data and were longer than all short intervals. Note that this is only a small fraction of the *number* of ISIs in the spike train (typically less than about 10%), since a few long intervals consume the same amount of time as many short intervals. Data from 300–1950 msec were processed, avoiding the transient portion of the response and the end of the spike train where the final inter-spike interval is unknown. With the longest intervals neutralized, the peaks were pushed down to the level of the noise in the ICC (thin lines, Fig. 5.10). Thus, 90% of the action potentials are serving to set a mean rate, and it is a few periods of long ISIs

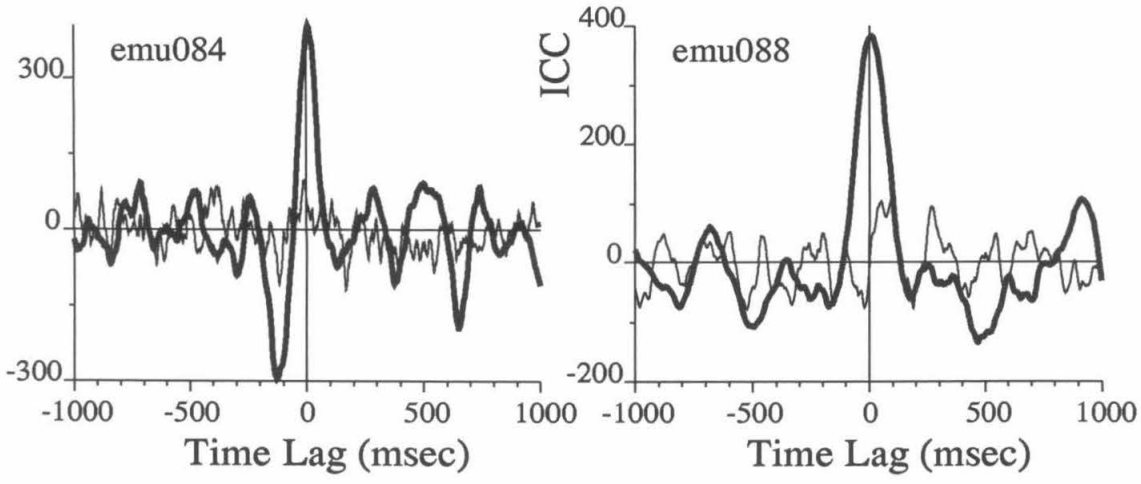

Figure 5.10: Peaks in the interval cross-correlations for two pairs of neurons (thick lines) are extinguished when the longest ISIs are removed (thin lines), leaving 70% of the data (over 90% of ISIs) intact. This demonstrates that long ISIs can provide a dominant source of correlation.

that dominate the ICC peaks.

5.3 Implications of Gaps and Bursts

We have identified a symmetry in the time course of the responses to preferred and null direction motions in area MT neurons which occurs in the absence of stimulus pattern induced temporal modulation. The occurrence of bursts during the coherent null direction motion is the less robust phenomenon but might be easier to explain than the gaps. One hypothesis is that the bursts result from the activation of a transient low threshold Ca^{++} current, I_T, from which inactivation has been removed by the strong inhibition of opponent inputs. This is known to elicit burst discharge.

We know of no internal mechanism that would cause the gaps, and therefore hypothesize that they may be the result of inhibition applied externally. The gaps seem to be a veto on the brisk firing induced by preferred motion; they appear in the midst

of high firing with no apparent gradual slowing down of the firing rate that is characteristic of adaptation. The gaps are unlikely to be caused by cessation of activity in the excitatory fibers coming from V1 because a motion energy model applied to the sparse dot stimulus shows that the excitation to individual V1 cells is transient in nature (see Chapter 6). We propose that the gaps arise from a computation organic to area MT. In Chapter 6, we show that the dynamic dot stimulus activates both null and preferred direction motion energy units, thus providing a possible drive for inhibition to the neuron during its preferred direction stimuli.

The observation that long ISIs, i.e., gaps, are correlated in 55% of pairs of neurons simultaneously isolated from a single electrode is also consistent with a common source of inhibition acting on both cells. A number of studies have proposed that inhibition plays an important role in giving MT cells their tuning properties (Mikami et al., 1986; Mikami, 1992). If inhibition causes the gaps, it may be from a different source than that proposed by Snowden et al. (1991) and Qian et al. (1994) if it comes from axoaxonic or basket cells making synapses onto the axon initial segment or soma, since this is unlikely to come from individual motion energy subunits. It has been reported recently that axoaxonic cells do show adaptation (Han, 1994), and this adaptation could be responsible for the limited time course of the gaps. Inhibition from large basket cells is believed to be non-specific for orientation in primary visual cortex based on anatomical evidence (Kisvarday and Eysel, 1993), so it would be interesting to see if the correlation between gaps varies with direction preference between pairs of MT neurons.

Neither the gaps nor the bursts display traits that are indicative of a stimulus artifact. They appear in data recorded over the course of 8 years for stimuli created by different software and displayed on different types of CRTs. The diversity of their time course and time of occurrence is characteristic of natural phenomena in cortex. However, it is necessary to find gaps in the response to other moving patterns if these

observations are to be taken seriously.

The monkey was awake and maintaining fixation during these experiments; therefore, it is worth questioning whether the gaps could be created by eyelid blink or by small saccades within the fixation window. It may be unlikely that the animal would blink during the stimulus presentation and risk breaking fixation. Also, variability of the timing of the gaps from neuron-to-neuron is inconsistent with eye blink; it is unlikely that the animal would consistently blink earlier for one neuron than for another. If the animal does blink, there is no obvious advantage to suppressing the response in area MT to stimulus motion during the blink. However, there is evidence that eye blinks are accompanied by a suppression of low spatial frequency vision which is of *neuronal* origin (Volkmann et al., 1982). It has been hypothesized that blink suppression and saccadic visual suppression may operate through the same pathways (Ridder and Tomlinson, 1993), and there is a clear advantage to suppressing responses to motion induced by saccadic eye movements. Although eye blink seems to be an unlikely cause for the gaps, it would not make the observation less interesting but would change some of our interpretations. Small saccades are more likely to account for the gaps than eyelid blink, and this possibility cannot be ruled out with the current data. Data from anesthetized monkeys, unable to blink or make saccades, could settle this issue.

We are pursuing further analysis to compare the correlation induced by the gaps to the correlation which causes narrow peaks, typically less than 10 msec wide, in standard cross-correlograms (for example, see Fig 6.18) that are often attributed to excitatory input (Ts'o and Gilbert, 1988; Kreiter and Singer, 1992). Both are sources of correlated spike count between pairs of cortical neurons. Correlation in spike count is an important factor that can limit the useful pool-size of neuronal ensembles (Zohary et al., 1994; Gawne and Richmond, 1993). Also, the long ISIs can provide an explanation for the power-law increase in variance of spike count with mean firing rate

that has often been reported in monkey (Vogels et al., 1989; Snowden et al., 1992). We have demonstrated this with a simple statistical model (unpublished results).

5.4 Response Variability

The high variability of spike trains has been noted in a number of recent studies (Softky and Koch, 1993; Shadlen and Newsome, 1994; Bell et al., 1995) and many older studies, some of which will be noted below. In this section, we briefly consider the variability of the MT spike trains in terms of spike count and inter-spike interval. Now that we have recognized the lack of stimulus-locked temporal modulation in the response to coherent motion stimuli ($c = 1$), we thought it is worth reconsidering variability in these terms.

We show that a statistical model motivated by the presence of long gaps in the spike trains can account for a power-law relationship between the mean and variance of spike count. We also consider implications of the small difference between the variability of $c = 1$ and $c < 1$ data in terms of coefficient of variation.

5.4.1 Variability of Spike Count

When spike count in a fixed time interval is used as the response measure for a neuron, the plot of variance versus mean response has revealed slopes on a log-log plot that are typically larger than 1.0 (Tolhurst et al., 1981; Dean, 1981; Bradley et at., 1987). Vogels et al. (1989) were the first to estimate response variability of striate neurons in behaving monkeys (using stationary square wave gratings). They reported a power of 1.11 and a variance that was 1.9 times the mean. Snowden et al. (1992) were the first to do a similar computation in an extrastriate area of a behaving monkey (using moving random dot patterns). They reported a power of 1.21 in V1 (intercept

1.08) and a power of 1.10 in MT (intercept 1.37). Softky and Koch (1993), in a different analysis of the MT database studied here, reported an exponent of 5/4 in the relationship between variance and mean of spike count.

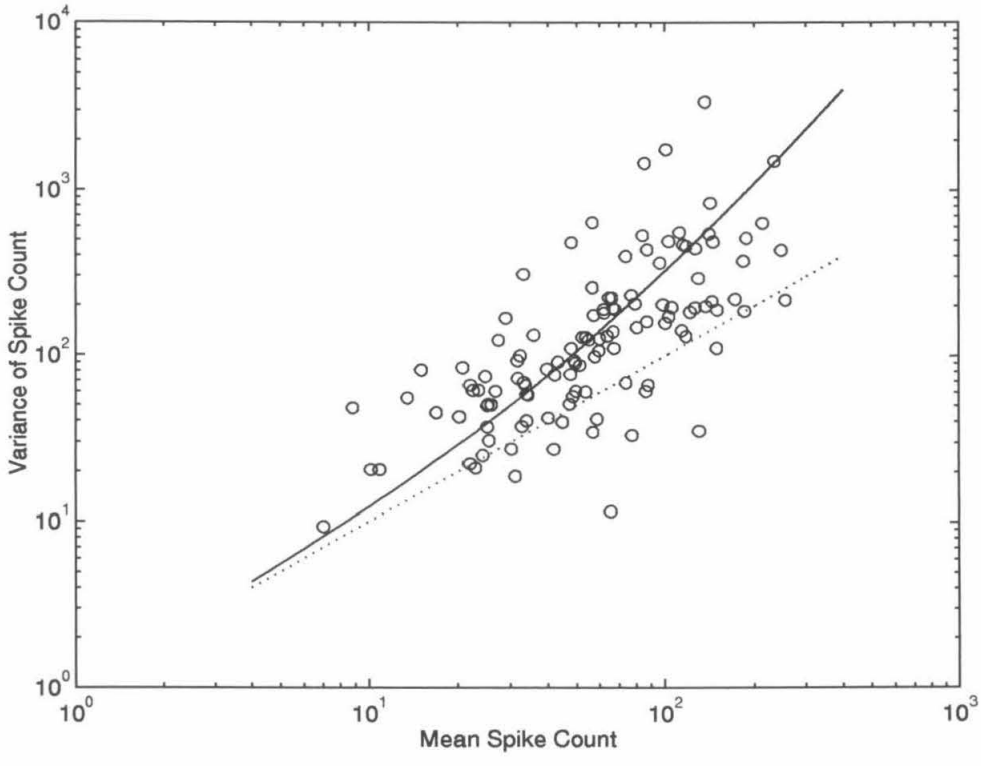

Figure 5.11: Response variance plotted agains response mean for spike count during 400–2000 msec. The dotted line shows the prediction from a Poisson process, while the solid line shows a prediction based on Eqn. 5.5 with $r = 0.15$.

According to numerous studies, this relationship appears to hold irrespective of what stimulus is used to achieve the mean rate (Dean, 1981; Tolhurst et al., 1981; Vogels et al. 1989). Vogels et al. (1989) stated, "This finding suggests that the relationship between response strength and response variance is a very general one

and is unrelated to the stimulus variable responsible for determining the response strength." Snowden et al. (1992) reiterated this statement.

We have found that a statistical model using compound distributions can account for a higher power in the relationship between variance and mean spike count. The motivation for this model comes from the observation of long gaps reported at the beginning of this chapter. A formal description is given below.

Let the spike count SC on a trial be the sum of the spike counts X_i from N independent time intervals of length T.

$$SC = X_1 + X_2 + \ldots + X_N. \tag{5.1}$$

Let the number of spikes X_i be Poisson distributed, with mean rate λ. Then,

$$\mathrm{E}(X) = \mathrm{VAR}(X) = \lambda T. \tag{5.2}$$

Now let N be distributed with mean μ and standard deviation σ. From the result in the Appendix of Chapter 7,

$$\mathrm{E}(SC) = \mathrm{E}(N)\mathrm{E}(X) = \mu\lambda T, \tag{5.3}$$

$$\mathrm{VAR}(SC) = \mathrm{E}(N)\mathrm{VAR}(X) + \mathrm{VAR}(N)\mathrm{E}^2(X) = \mu\lambda T + \sigma^2(\lambda T)^2. \tag{5.4}$$

Defining $M = \mu\lambda T$, the mean spike count, and $r = \sigma/\mu$, the coefficient of variation of the number of intervals (or length of time, as $T \rightarrow 0$) during which the Poisson process operates, then

$$\mathrm{VAR}(SC) = M + r^2M^2. \tag{5.5}$$

Choosing $r = 0.15$, a reasonably small value, produces the solid line in Fig. 5.11. The dotted line shows the result expected for a Poisson process (mean = variance). The

points show data from all MT cells with more than 5 trials at $c = 1$. Therefore, we conclude that the presence of long gaps, as reported above, can cause the relationship between variance and mean spike count to have an apparent power-law dependence, with exponent between 1 and 2. Usher et al. (1993) have previously accounted for variability in terms of a network model using feedback and inhibition that generate spike trains having statistics indicative of fractal processes.

5.4.2 Variability of Inter-spike Interval

Another commonly used method for characterizing variability is through the computation of C_V, the coefficient of variation (the ratio of the variance to the mean) of the inter-spike interval distribution (see references in Softky and Koch, 1993). Since there is little, if any, stimulus-locked temporal modulation for $c = 1$ stimuli, we wondered if the signal arriving at the spike generating zone might be less variable for $c = 1$ than for $c = 0$. For $c = 0$ stimulation, there must clearly be large fluctuations to produce the temporal modulation noted in Chapter 4. By analogy to the results of Mainen and Sejnowski (1995), is it possible that the apparently constant input, $c = 1$, generates a more regular (although less stimulus-locked) spike train? Since C_V is known to vary with spike rate, we computed the C_V at a fixed spike rate for $c = 1$ stimuli and separately for all other stimuli, $c \leq 0.5$. Fig. 5.12 shows that there is only a small, but statistically significant, reduction in the C_V of the ISI for coherent stimulation, indicating that the spike trains are slightly less variable for coherent motion. This is somewhat at odds with the proposals of past studies which indicated that variability was a function of spike rate, and was independent of the stimulus used to obtain that spike rate. However, the difference in the C_V is small, 0.06 on average across all spike rates, and we will consider the implications of this in the final section of this chapter.

Our C_V values, computed from the sustained period of the response, range from 0.8

to 1.0 for $|c| < 1$ stimulation (Fig. 5.12), and are somewhat less at $c = 0$ stimulation. Thus, we find that MT neurons, even when stimulated with random input, produce spike trains that are on average slightly more regular than a Poisson process. This is most likely related to the refractory period, discussed in Chapter 7 and Chapter 8.

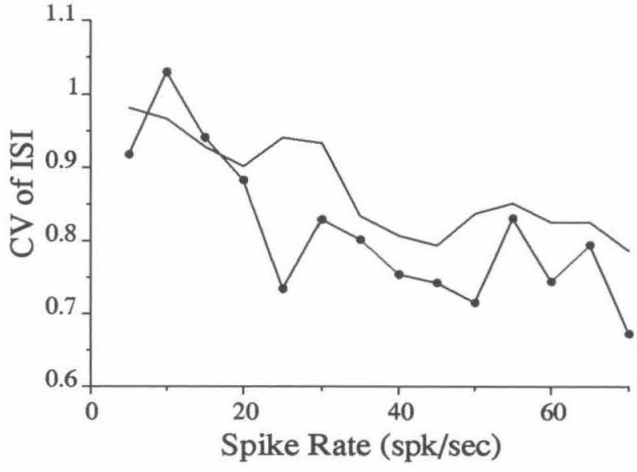

Figure 5.12: The variability of the ISIs, quantified by C_V, is plotted as a function of spike rate for coherent motion stimuli ($c = 1$, dots) and for non-coherent motion ($c \leq 0.5$). Each point is averaged over all cells. Bursts of spikes are eliminated by changing the spike trains to event trains as described in Chapter 7. The C_V is on average 0.06 less for coherent stimulation. This difference is statistically significant ($p < 0.005$), but small.

5.4.3 Unexplained Variability?

What sort of signal does the high-contrast coherent preferred motion stimulus supply to the spike initiating zone of the MT neuron? The first answer might be that the cell is saturated by its optimal stimulus, and should be receiving a constant, high level of excitation. There are two problems with this answer. First, cross-correlograms between pairs of neurons should be flat if either of the cells is saturated. Second, constant current injection to neocortical pyramidal cells *in vitro* produces a regular

train of action potentials with regular adaptation (although this is not true *in vivo*—see Holt et al., 1995). Neither of these conditions are satisfied here. However, a modulated, e.g. white noise, signal injected *in vitro* produces a spike train that is irregular but precisely locked to the injected signal (Mainen and Sejnowski, 1995). The apparent answer is that the MT cells receive a highly modulated signal, but one that is not locked to the visual stimulus. Narrow peaks in the cross-correlograms due to excitation is one method for identifying modulation in the input signals. The peaks reported here, caused by what might be common inhibition, may result from input supplied separately to the soma or initial segment of the axon and act as a veto on the signals originating in the dendrites. This is nevertheless still a source of modulated input to these MT cells. We cannot yet tell whether this modulation is an amplification of noise or whether it carries signals that are important for the functioning cortical network.

Chapter 6

Modeling Spike Trains from Area MT

This chapter describes the application of a motion energy model to the dynamic dot stimulus. We wanted to know whether the precise temporal modulation reported in Chapter 4 was consistent with the output of this class of computational model, which has been widely compared to electrophysiological data from both area MT and its V1 inputs (Heeger, 1987; Grzywacz and Yuille, 1990; Emerson et al., 1992; Qian et al., 1994; Nowlan and Sejnowski, 1994 and 1995). This model has led us to two major conclusions. First, the integration of the outputs of motion energy units can be made consistent with the *presence* of precise temporal modulation in the MT spike trains. Output from the model is shown, and the motion energy mechanism is analyzed in the frequency domain. Without knowing the exact stimulus patterns that were used in the original experiments, we cannot directly link the peaks in the PSTHs to the model output. Therefore, further experimentation is necessary to decide whether a motion energy model is an accurate description of the modulation.

The second major conclusion is an offshoot of the model. By varying the stimu-

lus speed, it became apparent that the stimulus was activating motion energy units tuned to motion in the opposite, or null, direction. An analysis of our numerical representation of the stimulus shows that its power spectrum is distributed in bands across the frequency domain. We compute an analytical expression for the power spectrum of the dynamic dot stimulus and show that it stimulates both preferred and null direction motion energy units within the optimal range of human psychophysical and V1 cell sensitivities, even at slow to moderate velocities for MT cells, i.e., 5–10°/sec. We discuss the difference between the power spectra of the signal dots and the noise dots and link this to the psychophysical observations from Chapter 3 and the stimulus-locked temporal modulation reported in Chapter 4.

6.1 Architecture of the Model

The model described here is similar to that of Adelson and Bergen (1985), and is therefore closely related to the Reichardt type of correlation motion detectors (Reichardt, 1961; van Santen and Sperling, 1984). Our model differs from that of Adelson and Bergen (1985) in that it is implemented in two spatial dimensions and has a final stage which integrates the output of opponent motion energy units across the spatial extent of an MT receptive field. For completeness, the next section describes the implementation of the motion energy filters. See Adelson and Bergen (1985) for a more thorough motivation. The second section describes the integration stage of our model.

6.1.1 Motion Energy Units

A motion energy unit consists of a pair of spatio-temporally oriented linear filters that are 90° out of phase, i.e., a quadrature pair. The output of the unit, called

motion energy by Adelson and Bergen (1985), is the sum of the squared output of the filters. The motion energy signal is directional and has the advantage that it does not depend on the phase of the stimulus, consistent with physiological results for directionally selective complex cells in V1 of macaque monkey (Emerson et al., 1992). One motion energy unit is encircled by the dotted line in Fig. 6.1.

The oriented linear filters of the motion energy units are often taken to be Gabor functions in space and time (Heeger, 1987; Grzywacz and Yuille, 1990). However, we followed the method of Adelson and Bergen (1985) and constructed quadrature pairs of oriented filters from sums and differences of separable linear filters composed of a Gabor function in space (two-dimensional in our case) and a *causal* temporal filter, which is mathematically less convenient but which is more biologically plausible. The construction of the four linear filters is diagramed at the top of Fig. 6.1, and the point along the flow of the diagram that corresponds to the convolution of the image with those filters is labeled by the filter names, f_{p1} and f_{p2} for the inputs to the preferred direction motion energy unit and f_{n1} and f_{n2} for the null direction unit. The equations describing the oriented filters and their separable components are given below. The temporal filters are characterized in the frequency domain because the frequencies which these filters pass will be related to the precision of the temporal modulation.

The construction of the oriented linear filters from two-dimensional spatial filters, g_1 and g_2, and from causal temporal filters, h_1 and h_2, is described by

$$f_{p1}(x,y,t) = g_1(x,y)h_1(t) + g_2(x,y)h_2(t) \tag{6.1}$$

$$f_{p2}(x,y,t) = g_1(x,y)h_2(t) - g_2(x,y)h_1(t) \tag{6.2}$$

$$f_{n1}(x,y,t) = g_1(x,y)h_1(t) - g_2(x,y)h_2(t) \tag{6.3}$$

$$f_{n2}(x,y,t) = g_1(x,y)h_2(t) + g_2(x,y)h_1(t), \tag{6.4}$$

where f_{p1} and f_{p2} are tuned to the preferred direction of motion, while f_{n1} and f_{n2}

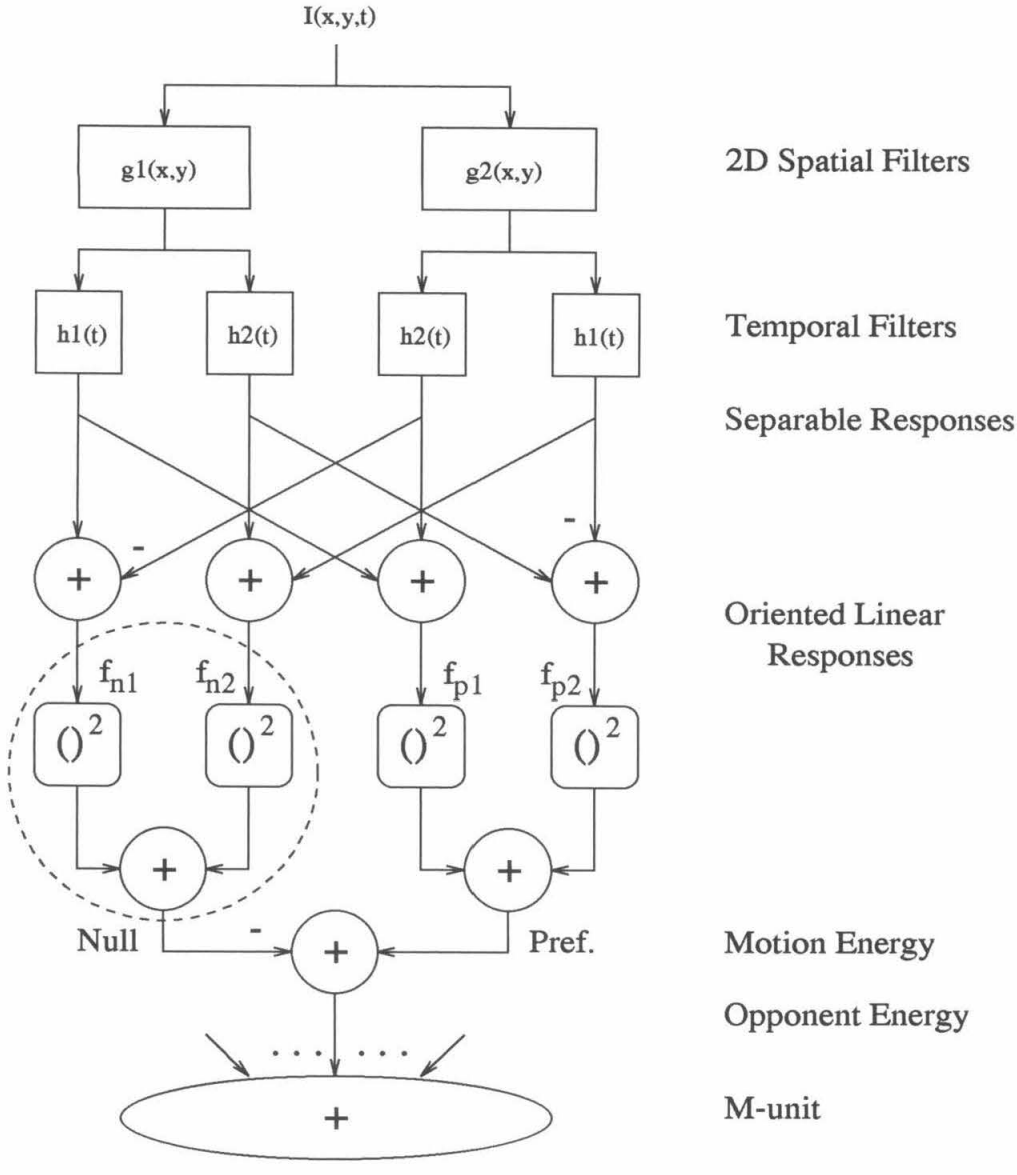

Figure 6.1: The structure of the motion energy model. One opponent energy MT sub-unit is shown contributing to our model "M-unit." The dotted circle indicates the null direction motion energy unit. The filters for the preferred and null motion energy units are constructed as prescribed by Adelson and Bergen (1985; Fig. 18b).

are tuned to the null direction (180° opposed to the preferred). The power spectrum of these three-dimensional filters is discussed in Section 6.4. The spatial filters are the even and odd Gabor functions,

$$g_1(x,y) = cos(2\pi f_s x)e^{-\frac{x^2+y^2}{\sigma_s^2}} \tag{6.5}$$

$$g_2(x,y) = sin(2\pi f_s x)e^{-\frac{x^2+y^2}{\sigma_s^2}}, \tag{6.6}$$

where f_s is the spatial frequency and σ_s is the spatial spread. Ideally, these parameters are set to match those of V1 neurons. The model that results from these equations will be tuned to upward motion. To achieve other orientations, the stimulus is rotated in space before the filters are applied. The spatial extent of the receptive field is a function of the preferred spatial frequency,

$$\sigma_s = \frac{1}{2f_s}. \tag{6.7}$$

This is consistent with reports that macaque V1 receptive fields contain typically 1–2 cycles of their optimum spatial frequency (Foster et al., 1985). The spatial filters $g_1(x,0)$ and $g_2(x,0)$ are shown in Fig. 6.2 (left) for the values $f_s = 1.0$ cycl/degr ($\sigma_s = 0.5$ degr).

Our model uses a causal temporal filter, unlike the more analytically tractable Gabor functions used by Heeger (1987) and Grzywacz and Yuille (1990). The temporal filter is derived from the difference of two multi-stage low-pass filters, one with $n+1$ stages and the other with $n+3$ stages. Each low-pass filter has a decaying exponential impulse response function (like that of an RC filter), so an m-stage filter has impulse response

$$a_m(t) = \frac{1}{\tau(m-1)!}\left(\frac{t}{\tau}\right)^{m-1} e^{-t/\tau} \qquad (t \geq 0). \tag{6.8}$$

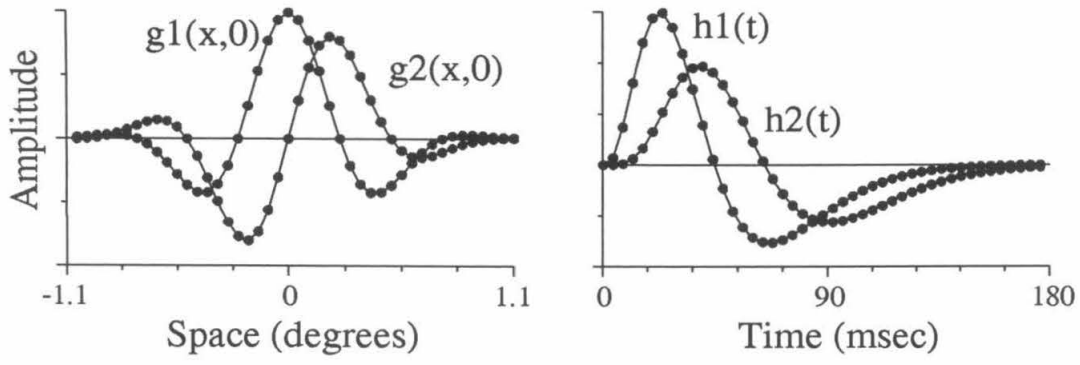

Figure 6.2: Spatial and temporal filter waveforms. The spatial filters are Gabor functions along the dimension shown, and simply Gaussian along the orthogonal dimension. The temporal filters are differences of multi-stage low-pass filters, used by Adelson and Bergen (1985).

It is interesting to note that this is the gamma probability density function that is used in Chapter 8 to model the refractory period for neuronal firing (set $a = m$, $b = 1/\tau$, and note that $\Gamma(m+1) = m!$ for integers). Both are the result of multi-fold convolution of the decaying exponential function. In the case of the refractory period, the convolution is related to summing independent random variables with the decaying exponential distribution, while in this case the convolution derives from cascading linear filters based on a simple RC circuit. The filter transfer function, $A_m(f)$, is

$$A_m(f) = (1 + i2\pi f\tau)^{-m}. \tag{6.9}$$

The temporal filter used by Adelson and Bergen (1985) is the difference between an $n+1$ and $n+3$ stage low-pass filter, as defined in Eqn. 6.8, with the definition $k = 1/\tau$, where τ is the time constant of the exponential decay for a single stage. This is the temporal filter of our model. Parameterized by n_i and k, the filter is

$$h_i(t) = \frac{a_{n_i+1}(t) - a_{n_i+3}(t)}{k} \tag{6.10}$$

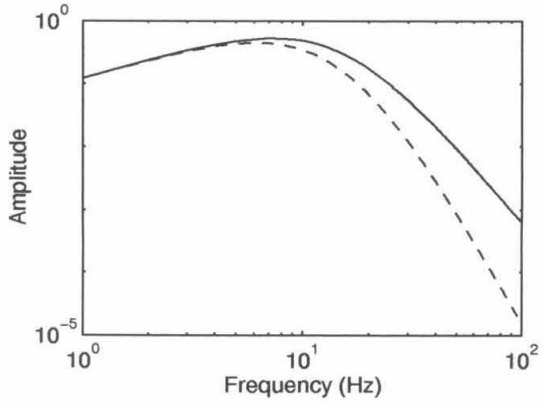

Figure 6.3: The frequency response of the temporal filters h_1 ($n = 3$, solid line) and h_2 ($n = 5$, dashed line). For both, $k = 100$.

$$= (kt)_i^n e^{-kt} \left[\frac{1}{n_i!} - \frac{(kt)^2}{(n_i+2)!} \right]. \tag{6.11}$$

The factor of $1/k$ keeps the equation consistent with Adelson and Bergen (1985) but makes no difference to the model. A similar temporal filter is used by Watson and Ahumada (1985) and Emerson et al. (1992), and they claim that it provides a good fit to human psychophysical data (reported by Robson, 1966) and to the temporal responses of V1 cells. The transfer function corresponding to Eqn. 6.11 is

$$H(f) = \frac{1}{k} \left[(1 + i2\pi f\tau)^{-(n+1)} - (1 + i2\pi f\tau)^{-(n+3)} \right]. \tag{6.12}$$

The temporal filters $h_1(t)$ and $h_2(t)$ are shown in Fig. 6.2 for the values $n_1 = 3$, $n_2 = 5$, and $k = 100$ ($\tau = 10$ msec). Fig. 6.3 shows the modulus of $H(f)$ for the same parameters. Scaling $h(t)$ by changing the value of k will cause an inverse scaling of $H(f)$, so the filters can be adjusted in this way to have cutoff frequencies in the range observed for V1 neurons, typically 4–16 Hz (Foster et al., 1985).

6.1.2 Integration of Motion Energy

The next stage of the model involves the integration of the outputs of many motion energy units. Our approach will be to average the outputs of opponent energy units across an area the size of an MT receptive field. This is described in detail below, but first a few other approaches are noted.

A number of models have combined motion energy outputs to create velocity tuned units with some similarities to MT neurons (Heeger, 1987; Grzywacz and Yuille, 1990; Nowlan and Sejnowski, 1994 and 1995). The models of Heeger (1987) and Grzwacz and Yuille (1990) begin with three-dimensional Gabor functions to extract motion energy. Heeger's model estimates the amount of energy across a local region of space for a small set of motion energy filters (12) arranged as a cylinder in the spatio-temporal frequency domain. To achieve velocity tuning, the strong assumption is made that the power spectrum of the input is flat. This allows the model to sample so few points in the frequency domain. The model of Grzywacz and Yuille (1990) does not make assumptions about the power spectrum of the input and computes a velocity estimate at each point in space, i.e., without spatial averaging, from a weighted average across a plane in the frequency domain (their "ridge strategy"). The resulting velocity tuned units have non-separable sensitivities in the frequency domain. Neither of these models account for the large receptive fields of MT cells. In Heeger's model, which performs spatial averaging, the local spatial integration region is much smaller than the extent of an MT receptive field, being only four times larger in area than the Gabor filters that represent V1 cells. MT receptive fields may be 50–100 times the area of V1 receptive fields at a similar eccentricity, since the MT receptive field diameters are about 10 times those of V1 cells at all eccentricities (Gattass and Gross, 1981; Albright and Desimone, 1987; Maunsell and Newsome, 1987). Grzywacz and Yuille (1990) make no explicit account for the size

of MT receptive fields.

The model of Nowlan and Sejnowski (1994, 1995) uses two parallel pathways to estimate not only the local velocities across spatiotemporal frequency channels but also the *validities* of those estimates. Their model of selective integration of signals derived from local motion energy is able to replicate the psychophysical and physiological findings of Britten et al. (1992) and does account for the integration of information across realistic MT receptive field sizes. It would be interesting to test how their model would account for the temporal modulation reported here.

In a fourth model, at the level of MT sub-units rather than velocity tuned units, MT responses to transparent and non-transparent stimuli were accounted for by integrating motion energy locally in an opponent, or "suppressive," stage (Qian et al., 1994). In this model, the opponent stage of Adelson and Bergen's motion energy model constituted an MT sub-unit. The output of such a sub-unit, shown at the bottom of Fig. 6.1, is the preferred minus the null direction motion energy. (We think this suppression might occur earlier than the inhibition that might be responsible for the gaps in the spike trains reported in Chapter 5.) Our "MT sub-unit" is similar to that of Qian et al., except we use a causal temporal filter in the motion energy stage, as given in Eqn. 6.11.

How an MT neuron is endowed with its direction and speed tuning is still uncertain, so we will not attempt to make a complete model of an MT cell. Our model unit, called an M-unit, is an artifice that will be compared to an MT neuron in terms of the data analyzed here and clues from the literature. The M-unit simply averages the outputs of opponent MT sub-units across the extent of its circular receptive field. An M-unit is defined by the following parameters

1. Receptive field diameter, d_{RF}
2. Spatial frequency of the sub-units, f_s
3. Spatial orientation of the sub-units, θ

4. Temporal filters of the sub-units, k, n_1, n_2

One important difference between the M-unit and some models of MT neurons (for example, the velocity cells of Grzywacz and Yuille, 1990) is that the M-unit is not velocity tuned. The M-unit inherits the tuning of its opponent motion energy sub-units, and that tuning, in the frequency domain, consists of two volumes with roughly teardrop shaped contour surfaces aligned with the temporal frequency axis. (Projections of these are shown later in Section 6.4.) The M-unit, therefore, has approximately separable tuning within one frequency domain octant, i.e., the response is approximately the product of three one-dimensional functions, one along each frequency domain axis. The dependency of optimal velocity on spatial frequency for a separable filter is demonstrated in Fig. 6.4. Of the few area MT cells that have been recorded with varying spatial and temporal frequency sinusoids, most have tuning that would be better described as separable than as aligned on an iso-velocity contour (J. Anthony Movshon, unpublished data), implying that the preferred velocity *does* change as spatial frequency (or temporal frequency) changes.

The model includes an explicit representation of the population of sub-units which contribute to the M-unit. The sub-units are spaced σ_s degrees apart on a rectangular grid within the M-unit receptive field. This is shown in Fig. 6.5 for an M-unit with $d_{\mathrm{RF}} = 5°$ and $\sigma_s = 0.5°$. Adding more sub-units between the ones shown in the figure will not change the output of the M-unit. This follows from the sampling theorem and the approximate band-limiting that is performed by the spatial filters in the motion energy computation.

The final output of the M-unit is the half-wave rectified average of the sub-unit outputs. The output can be plotted as a probability of firing an action potential as a function of time. This can then be used as the time varying rate parameter of an inhomogeneous Poisson process to generate spike trains. Only here, at the generation

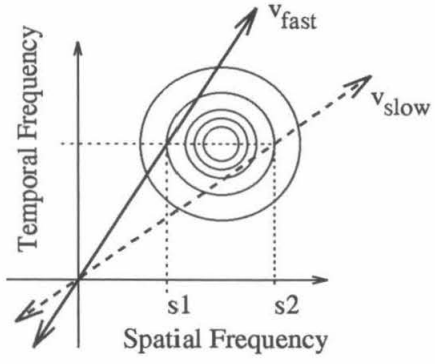

Figure 6.4: Motion energy filters are not tuned for velocity. The circles represent a contour plot of the tuning of a motion energy filter that is separable in spatial and temporal frequency—the preferred spatial and temporal frequencies are the coordinates of the center of the iso-amplitude contours. The solid line shows the preferred velocity of a pattern with spatial frequency concentrated at s_1, while the dotted line shows the *slower* preferred velocity for a higher spatial frequency, s_2. The velocity lines are chosen to intersect the spatial frequency lines at the maximum point with respect to the tuning surface of the filter. However, for an image with a flat spatial frequency spectrum, this motion energy filter will prefer the velocities that are constrained to pass through the center of the tuning contours.

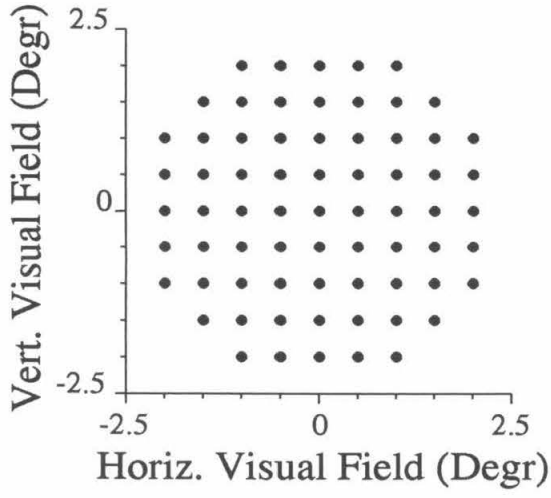

Figure 6.5: The locations of the centers of the 69 opponent motion energy sub-units which contribute input to one 5° diameter M-unit. The Gaussian receptive field (not shown) of each sub-unit has standard deviation $\sigma_s = 0.5°$. The sub-unit spatial frequency is $f_s = 1.0$ cycl/degr (see Eqn. 6.7).

of Poisson distributed action potentials, is noise introduced into the model. In the next section, we compare the output from the M-unit to MT data.

6.2 Results from the Model

Three points are illustrated using the model. First, the output of the motion energy stage, corresponding to directionally selective V1 complex cells, is highly modulated for both coherent, $c = 1$, and incoherent, $c = 0$, motion, as argued in Chapter 5. Second, the M-unit output is relatively unmodulated at $c = 1$ but is highly modulated at $c = 0$, showing a qualitative consistency with spike trains and PSTHs produced by the MT cells. Third, low frequency modulation remaining in the M-unit response to $c = 1$ motion is inconsistent with the MT data and implies that a normalization stage is necessary.

The input to the M-unit was a set of points in three dimensions representing the dynamic dot stimulus, reproduced exactly as described in Chapter 2. In particular, the 45 msec Δt and the distribution and density of dots were consistent with the actual experiments. (It turned out that exactly reproducing the stimulus led us to an important observation that is developed analytically in Section 6.4.) The parameters used for the M-unit are as follows: $d_{\mathrm{RF}} = 5.0°$, $f_s = 1.0$ cycl/degr, $\sigma_s = 0.5°$, $\theta = 90°$, $n_1 = 3$, $n_2 = 5$, and $k = 100$ ($\tau = 10$ msec).

A speed tuning curve and two direction tuning curves are shown in Fig. 6.6 averaged over 10 trials with randomly seeded stimuli. The preferred speed is 6°/sec and the preferred direction is, as defined, $\theta = 90°$. The direction preference is reversed at 16°/sec (thick line, right panel, Fig. 6.6). This is explained in Section 6.4. The remaining output in this section was computed at the preferred velocity, 6°/sec, unless otherwise noted.

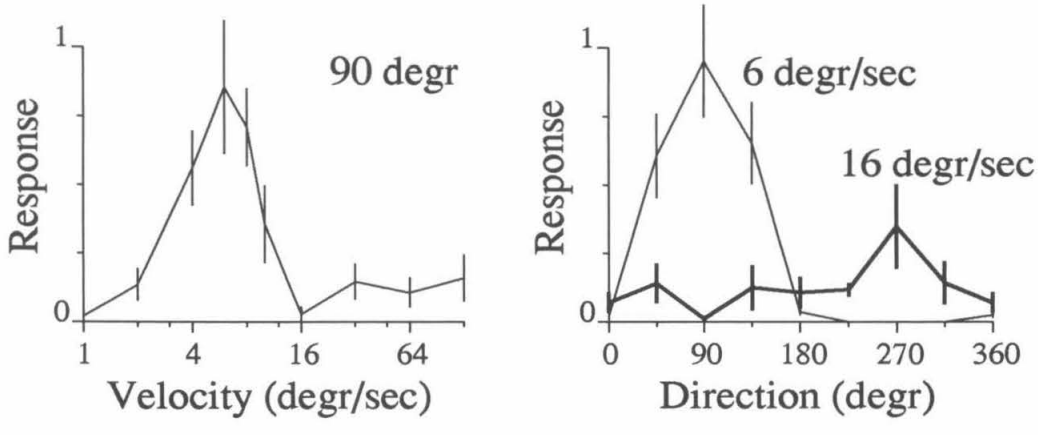

Figure 6.6: The M-unit velocity tuning (left) for $c = 1$ stimuli and the direction tuning (right) at the preferred velocity (6°/sec) and a higher velocity (16°/sec, thick line). The reversal of preferred direction at 16°/sec can be explained from the filter profiles in the frequency domain shown in Fig. 6.14. Ten trials with differently seeded stimuli were used to create these curves. Error bars show the standard deviation. The response is given in arbitrary units proportional to spike rate.

We argued in Chapter 5 that V1 cells were likely to have modulated responses to

the dynamic dot stimulus because of its low dot density. The V1 response is taken to be the output of a (non-opponent) motion energy unit. Fig. 6.7 shows a typical model V1 response to the dynamic dot stimulus at $c = 0$ and at $c = 1$. Both responses are modulated, but the modulation at $c = 1$ has more low frequency components due to the time it takes the coherently moving pattern to traverse the V1 receptive field. These responses are consistent with the sparse dot density.

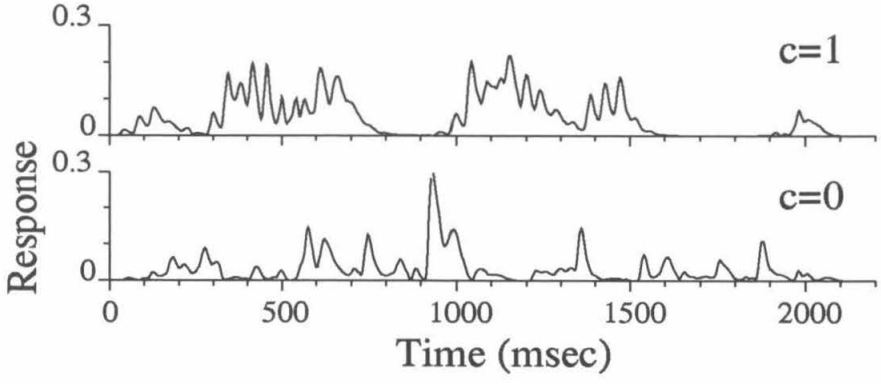

Figure 6.7: The normalized response of a preferred direction motion energy unit (like a V1 complex cell) to a $c = 1$ stimulus (top) and a $c = 0$ stimulus (bottom). The V1 receptive fields are small relative to the dot density of the stimulus, so the responses are highly modulated even for $c = 1$ in spite of the quadrature filters.

Now we examine the temporal structure of the M-unit response. Spike trains and PSTHs are shown for the M-unit in Fig. 6.8 and for an MT neuron in Fig. 6.9. The M-unit spike trains are generated from the time varying firing probabilities shown at the bottom of Fig. 6.8 (thick line $c = 0$, thin line $c = 1$). Both the M-unit and the MT cell have 22.2 Hz oscillations due to the 45 msec Δt of the dot stimulus, although not all MT cells respond to the 22.2 Hz signal. Both the M-unit and the MT cell have higher frequency modulation in the response to $c = 0$ motion. Fig. 6.10 compares the power spectra of the PSTHs from the spike trains for $c = 0$ and $c = 1$ for the MT neuron (left) and the M-unit (right). Both the cell and the model have

less low frequency modulation in response to coherent motion, and both have a peak at 22.2 Hz.

The frequencies observed in the spectra and the modulation in the spikes trains is largely due to the temporal frequency cutoff of $h_1(t)$ and $h_2(t)$. The filters used have higher than average cutoffs for V1 cells, but are not unrealistic. However, the half-wave rectification is also responsible for producing high frequencies. In the next section, we will compute the effect of the quadrature energy computation on the frequency spectrum of the signal and argue that the model can be made quantitatively consistent with precision values observed in the data in Chapter 4. Therefore, quantitative analysis of the temporal structure of the model output will not be pursued here.

The final point of this section is related to spike count rather than temporal structure. Although the output of the model is easily made consistent with the stimulus-locked temporal modulation, it does not match well the variability of the mean spike count averaged over the entire trial, particularly for coherent motion. Fig. 6.11 shows that the variance-to-mean ratio for the spike count during the sustained period of the trial (400–2000 msec, here) for MT neurons is roughly constant at 2 across all coherence levels (Fig. 6.11, right). The output of the model in terms of spike count, however, becomes variable for coherent stimuli (Fig. 6.11, left). The increase in variance-to-mean ratio at $c = \pm 1$ is more dramatic for a non-optimal velocity stimulus (filled circles) than for the optimal velocity stimulus (triangles, 10°/sec). The exaggerated variance for coherent motion is due to the relatively fixed spatial structure of the highly coherent stimulus, and its interaction with the locality of the motion energy filters in the frequency domain. For $c = 1$, the distribution of power in the spatial frequency domain changes slowly as the stimulus translates across the receptive field. In contrast, for $c = 0$ stimuli, the spatial structure is completely uncorrelated after 45 msec, since all dots have been replotted randomly in that time.

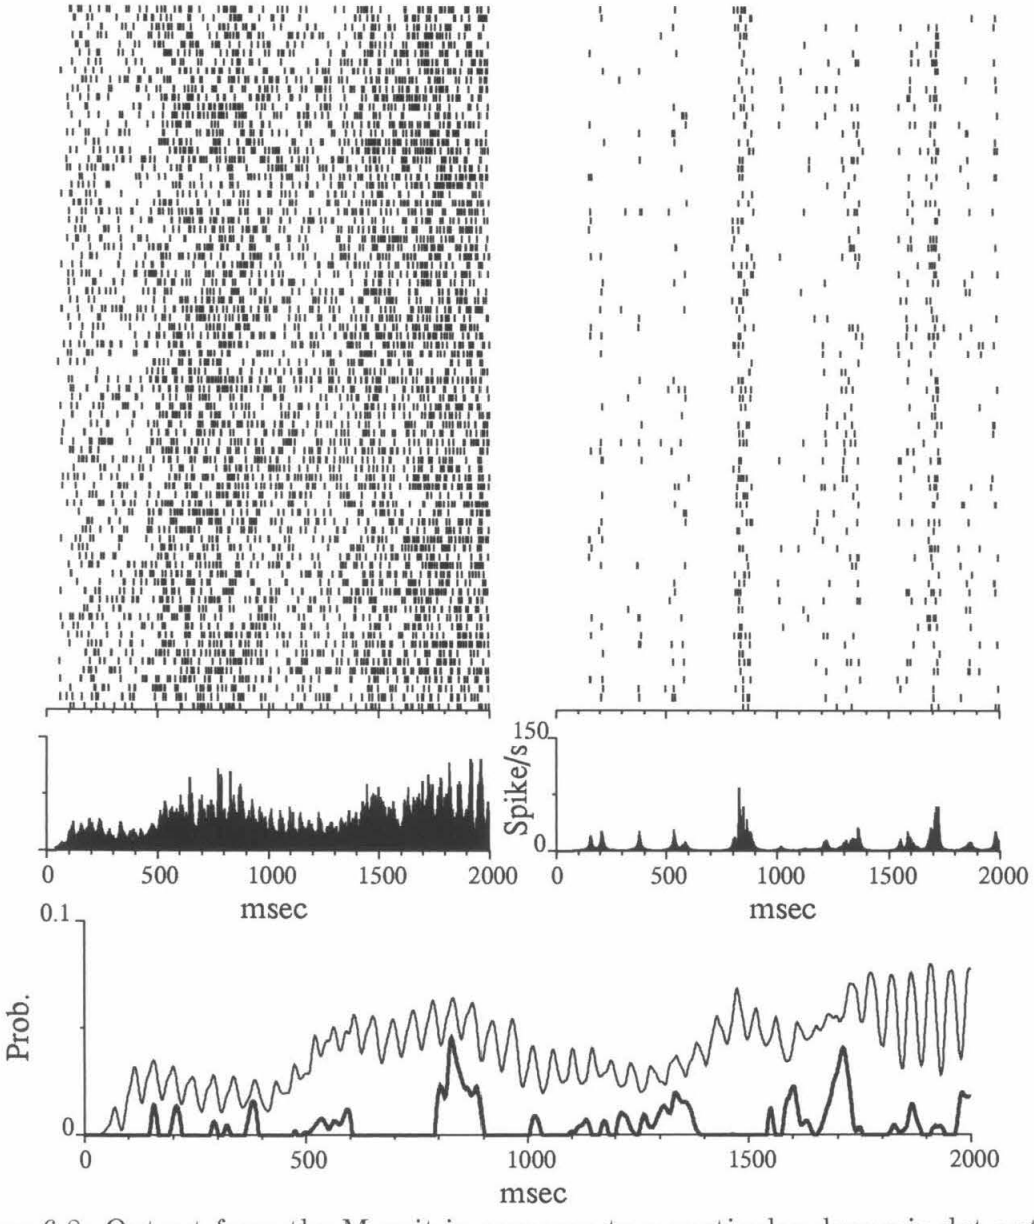

Figure 6.8: Output from the M-unit in response to a particular dynamic dot pattern at $c = 1$ and another at $c = 0$. Spike trains and PSTH for $c = 1$ (left) and $c = 0$ (right) were generated from the two curves at the bottom, which are the output of the M-unit for the $c = 1$ stimulus (thin line) and the $c = 0$ stimulus (thick line).

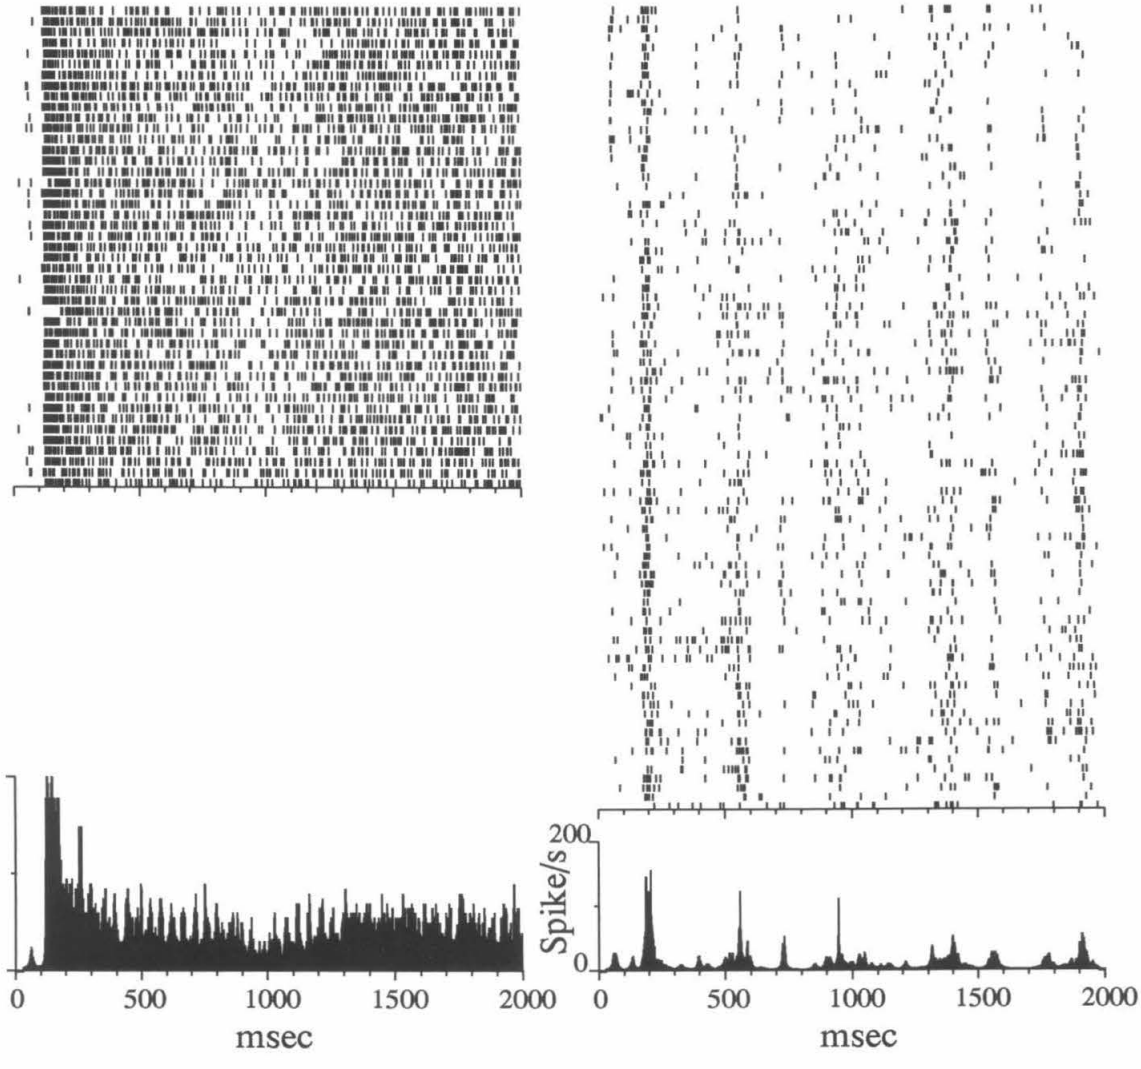

Figure 6.9: Data from an MT neuron (j025) for comparison against the model output in Fig. 6.8. Spike trains are shown in response to a particular dynamic dot pattern at $c = 1$ (left) and a pattern at $c = 0$ (right).

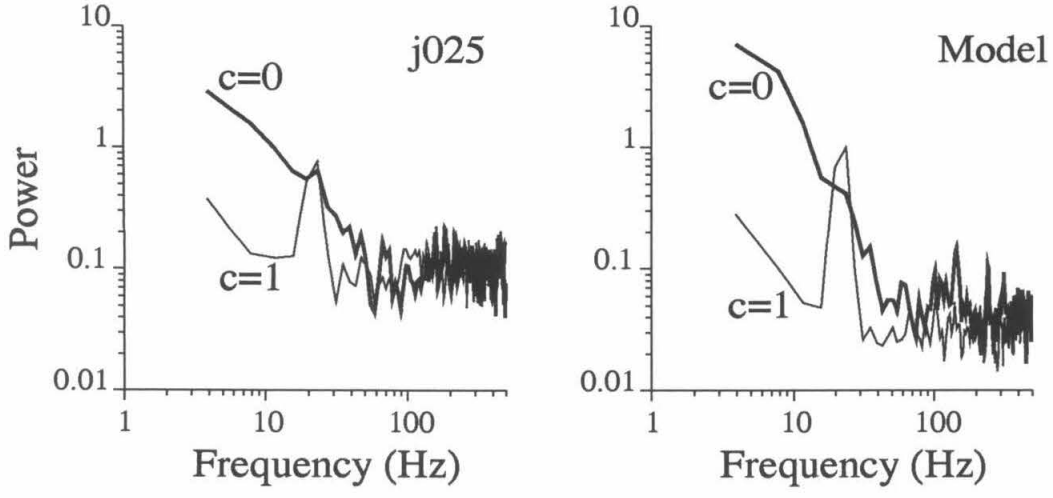

Figure 6.10: Power spectrum of PSTHs from the model (Fig. 6.8) and from the MT neuron j025 (Fig. 6.9). Both have a peak at 22.2 Hz for coherent motion, and both have more modulation at low frequencies, i.e., less than 60 Hz, for $c = 0$ than for $c = 1$.

The exaggerated variance is most likely the result of the narrow region of spatio-temporal frequency domain that is measured by an M-unit. This is likely because the variance in the number of dots from stimulus to stimulus is nearly the same regardless of the coherence, so the variance of the total power in the frequency domain also must be the same (by Parseval's theorem). This actually argues in favor of models that are velocity tuned, since these models compare the distribution of motion energy across the spectrum rather than measuring an absolute amount locally (Heeger, 1987; Grzywacz and Yuille, 1990).

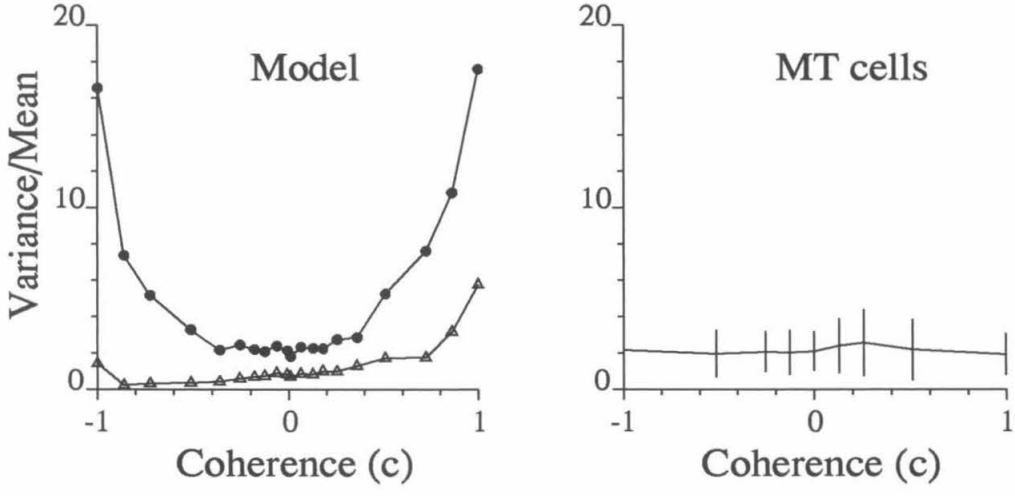

Figure 6.11: Variance-to-mean ratios for spike count in the model (left) and the MT data (right). At $c = \pm 1$ the variance-to-mean ratio increases for the model but not for the neuronal data. The model was tested at an optimal speed, 6°/sec (triangles), and a non-optimal speed, 10°/sec (filled circles). For the neuronal data, the parameters of the dynamic dot stimuli were optimized by the experimenter. Error bars show standard deviation for the MT data. Typically, data from 130 MT cells were included at each point (right).

6.3 Effect of Quadrature Energy on the Power Spectrum

This section presents an argument on analytical grounds that precise and high frequency stimulus-locked temporal modulation can be achieved using the motion energy model, and the computational model does not need to be tested further simply to make this point. This argument consists of an examination of the effect of the quadrature energy mechanism on the power spectrum of its input followed by an appeal to the results from Chapter 4 linking sinusoidal modulation of firing rate to temporal precision in terms of the jitter measure.

Consider a one-dimensional signal in time, $f(t)$, which is passed through a filter with impulse response $h(t)$ and, in parallel, through the quadrature filter of $h(t)$. The quadrature filter, $h_q(t)$, is obtained by computing the Hilbert transform of $h(t)$, i.e., convolving $h(t)$ with $-1/(\pi x)$. In Fourier space, this amounts to multiplying by $i\text{sgn}\omega$.

If the output of the filter $h(t)$ applied to the input $f(t)$ is called $s(t)$,

$$s(t) = f(t) * h(t), \tag{6.13}$$

then the quadrature energy output, $q(t)$, is

$$q(t) = [s(t)]^2 + \left[s(t) * \left(-\frac{1}{\pi t}\right)\right]^2. \tag{6.14}$$

In the frequency domain, this is

$$Q(\omega) = S(\omega) * S(\omega) + [S(\omega)i\text{sgn}\omega] * [S(\omega)i\text{sgn}\omega], \tag{6.15}$$

which may be rewritten as the integral

$$\int_{-\infty}^{\infty} S(u)S(\omega - u)\left[1 - \mathrm{sgn}u\mathrm{sgn}(\omega - u)\right] du. \tag{6.16}$$

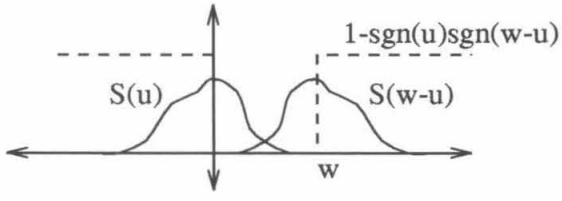

Figure 6.12: The effect of quadrature energy in the frequency domain. The factor involving "sgn" (dotted line), which comes from the Hilbert transform, prevents low frequencies from contributing to higher frequencies in the convolution integral. Auto-convolution in the frequency domain is the result of squaring in time.

This integral represents a modified auto-convolution of the Fourier transform of $s(t)$. The factor containing the sign function suppresses the auto-convolution within the window from zero to ω. This is shown graphically in Fig. 6.12. The nature of this windowing assures that frequencies at ω are not created by the product of two lower frequencies because only the portion of $S(u)$ in Fig. 6.12 that *exceeds* ω (and the portion of $S(\omega - u)$ that is more negative than $-\omega$) can contribute to the amplitude at ω in the resulting Fourier transform, $Q(\omega)$. In addition to introducing no new high frequency components, the operation creates a DC response that is equal to the auto-convolution of $S(\omega)$ evaluated at zero, which is the integral of $S^2(\omega)$ (since $S(\omega)$ is hermitian, $S(\omega) = S^*(-\omega)$, given that $s(t)$ is real and asymmetrical, see Bracewell, 1978).

This analysis demonstrates that the quadrature energy mechanism will not shift power to higher frequencies than those originally passed by the temporal filter un-

derlying the mechanism. Because the operation is similar to convolution, power is spread out. For example, a band-pass spectrum will become low-pass, and a low-pass spectrum will become more concentrated around the origin but will still contain power out to its original cutoff frequency. Therefore, frequencies that are passed by the temporal filter will be present in the output, and when converted to spike trains, can cause high frequency, precise temporal modulation. Fig. 4.8 in Section 4.3 shows that simply scaling the amplitude of a truncated signal can produce higher temporal precision. Truncation alone introduces higher harmonics; the spectrum of a truncated sine wave contains odd harmonics. Therefore, we conclude that the normalization and truncation of the signal can be as important as the transfer function of the temporal filter in determining the frequencies in the power spectrum of the output of the M-unit.

6.4 The Dynamic Dot Stimulus in the Frequency Domain

A frequency domain analysis of the random dot stimulus reveals that it spreads power across frequency space, and the spreading becomes more homogeneous as velocity increases. By modeling the details of the stimulus generation in both space and time, we extend the analysis which was first performed by Britten et al. (1993), who did not take into account the stroboscopic nature of the stimulus.

It is known that translation over time of a three-dimensional image results in a three-dimensional Fourier transform which falls along a plane in the frequency domain (Watson and Ahumada, 1983). It is also known that stroboscopic motion is represented in the frequency domain by a set of equally spaced parallel planes which includes the plane of the corresponding smooth motion. The slope of the plane is the

velocity of the motion.

The implementation of the dynamic dot stimulus, as described in Chapter 2, is not strictly stroboscopic, in the sense that samples are not taken at the same time across all space. Therefore, the analytical model of the stimulus used here will consist of two components. The signal component consists of pairs of δ-functions in space and time for which Δx and Δy are specified by the stimulus velocity, given that Δt is fixed at 45 msec, as in the experiments. The signal component comprises these pairs placed randomly in space and time, i.e., in accordance with Poisson statistics. The noise component consists of single δ-functions placed randomly in space and time. A stimulus with N dots at coherence c has cN dots in the signal component and $(1-c)N$ dots in the noise component.

Each pair of dots in the signal component can be described as a pair of δ-functions,

$$\phi(x,y,t) = \delta(x,y,t) + \delta(x+\Delta x, y+\Delta y, t+\Delta t). \tag{6.17}$$

The Fourier transform, $\Phi(u,v,w)$, of $\phi(x,y,t)$ is

$$\begin{aligned}
\Phi(\omega_x,\omega_y,\omega_t) &= \int\int\int [\delta(x,y,t) + \delta(x-\Delta x, y-\Delta y, t-\Delta t)] \\
&\quad e^{-i(\omega_x x+\omega_y y+\omega_t t)}\, dx\, dy\, dt && (6.18)\\
&= 1 + e^{-i(\omega_x \Delta x+\omega_y \Delta y+\omega_t \Delta t)} && (6.19)\\
&= 1 + \cos(\omega_x\Delta x + \omega_y\Delta y + \omega_t\Delta t) \\
&\quad -i\sin(\omega_x\Delta x + \omega_y\Delta y + \omega_t\Delta t). && (6.20)
\end{aligned}$$

The power spectrum is given by the square of the modulus,

$$|\Phi(\omega_x,\omega_y,\omega_t))|^2 = 2 + 2\cos(\omega_x\Delta x + \omega_y\Delta y + \omega_t\Delta t). \tag{6.21}$$

This is the power spectrum for one pair of dots. For low coherence stimuli, most of the signal component consists of pairs of dots. Assuming $c \geq 0$, the probability of a dot being refreshed once only, making a signal pair, is $c(1-c)$, while the probability of a dot being refreshed at least once, being involved in a signal of any length, is c. Therefore, the fraction of the signal dots that are in the form of pairs is

$$\Pr\{pair\} = \frac{c(1-c)}{c} = 1 - c. \tag{6.22}$$

Thus, at $c = 0.1$, 90% of the signal dots take the form of pairs. The pairs of dots have essentially random placement in space and time, so the signal component can be modeled as Poisson shot noise in three-space, convolved by $\phi(t)$. Poisson shot noise has a flat power spectrum (Champeney, 1973), each pair of signal dots has a cosine spectrum (Eqn. 6.21), and convolving in space and time corresponds to multiplication in the frequency domain. Therefore, the mean of the spectrum of the stimulus at low coherence is approximated by a cosine. At high coherence, the signal component is approximated as regularly sampled motion of many moving dots at random phases. The spectrum in this case is a series of planes, as expected for phase-locked stroboscopic motion. The peaks in the spectrum, however, fall along the same planes for both high and low coherence.

To compare this result to the two-dimensional projections in Fig 6.13 and Fig 6.14, we assume $\Delta x = 0$ since motion was in the upward direction. The maximum values of Eqn. 6.21 occur when the argument of the cosine is an integer multiple of 2π; therefore, when

$$\omega_t = -\frac{\Delta y}{\Delta t}\omega_y + \frac{k}{\Delta t}, \qquad (k = 0, \pm 1, \pm 2 \ldots). \tag{6.23}$$

The stimulus had $\Delta t = 0.045$ sec, and considering the case for $k = 1$, there should

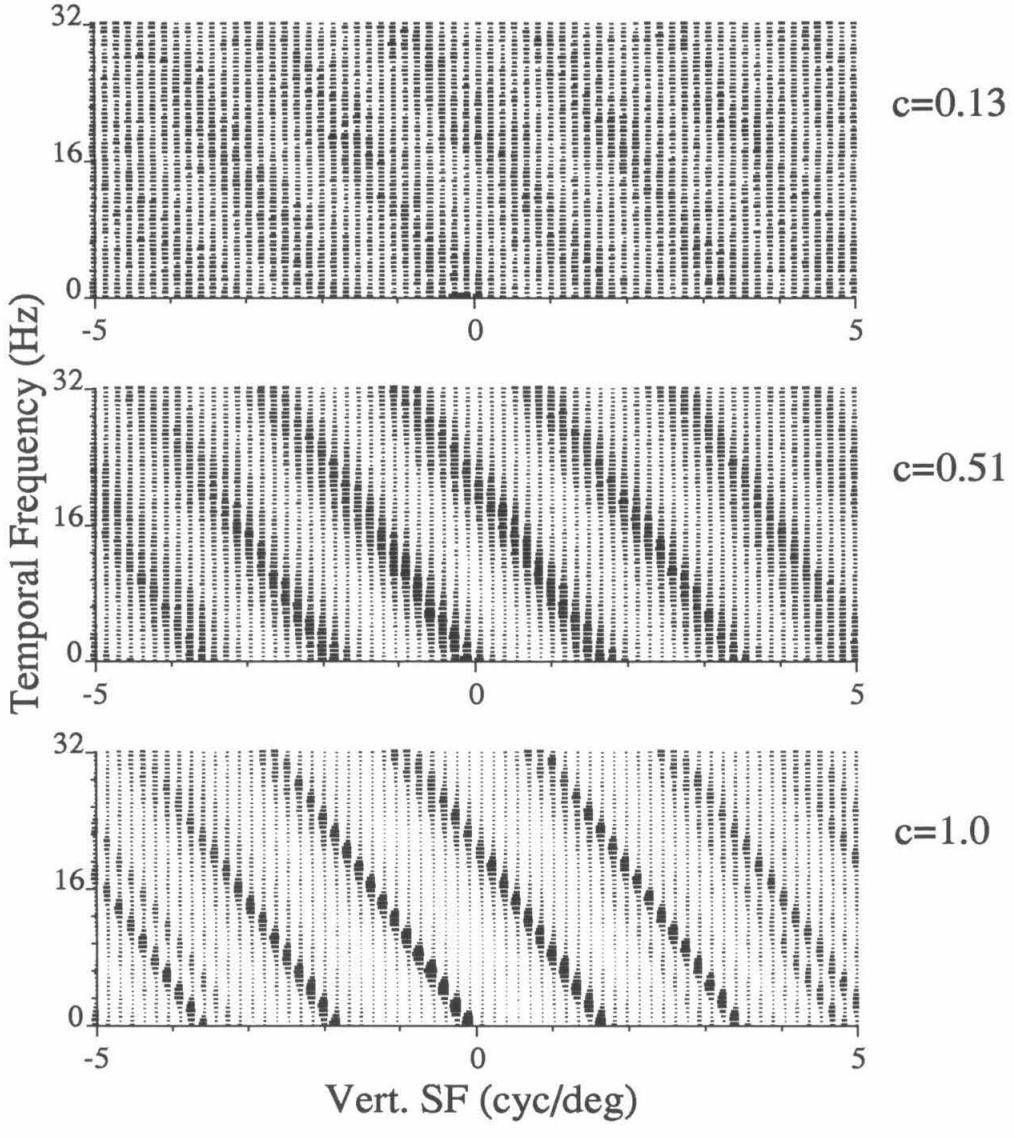

Figure 6.13: Power spectrum of the numerical representation of the dynamic dot stimulus for three coherence levels. The parallel planes for the $c = 1$ stimulus (C) blurs into the cosine shape described by Eqn. 6.21 for lower coherence values (A and B).

be a peak in the spectrum along the line

$$\omega_t = -v_y\omega_y + 22.2, \tag{6.24}$$

where $v_y = \Delta y/\Delta t$ is the speed. This equation predicts a peak along a line with y-intercept 22.2 Hz and having slope equal to the negative of speed. This matches the projections in Fig. 6.13, which were computed from our numerical representation of a 10°/sec stimulus at $c = 0.13$, $c = 0.51$, and $c = 1.0$. In the figure, the parallel lines (projections of planes) for $c = 1.0$ are blurred into the cosine form at low ($c = 0.13$) coherence. Fig. 6.14 shows power spectra for $c = 1$ stimuli at 5°/sec (C) and 10°/sec (D). As the speed increases, the slope becomes steeper and the peaks get closer together. At 10°/sec, a number of peaks along lines other than the line through the origin traverse spatio-temporal frequencies within the range of V1 sensitivity.

The spectra here appear somewhat different than those reported by Britten et al. (1993, Fig. 10). It appears that they approximated their stimulus by moving the dots only one discrete step at a time in their numerical representation. This puts the first copy of the spectrum at the cutoff frequency which is at the edge of the plot. As the parallel planes spread into a cosine pattern for low coherence motion, the copies of the spectrum appear at the corners of their plots F and G. Their plots look like the plots here if one considers that their temporal frequency axis extends to ±11.1 Hz (given the sampling was set to 45 msec) and their spatial frequency axis extends to $\pm 11.1/v$ cycl/degr, where v is the velocity of the stimulus. It appears to be common practice to choose the discrete step size of the numerical simulation to be equal to the Δt of the stroboscopic motion; this appears to be the case in the Fig. 5 of Qian et al. (1994). In general, it would be desirable to use a numerical sampling rate *higher* than the stroboscopic rate so that the copies of the power spectrum appear below the

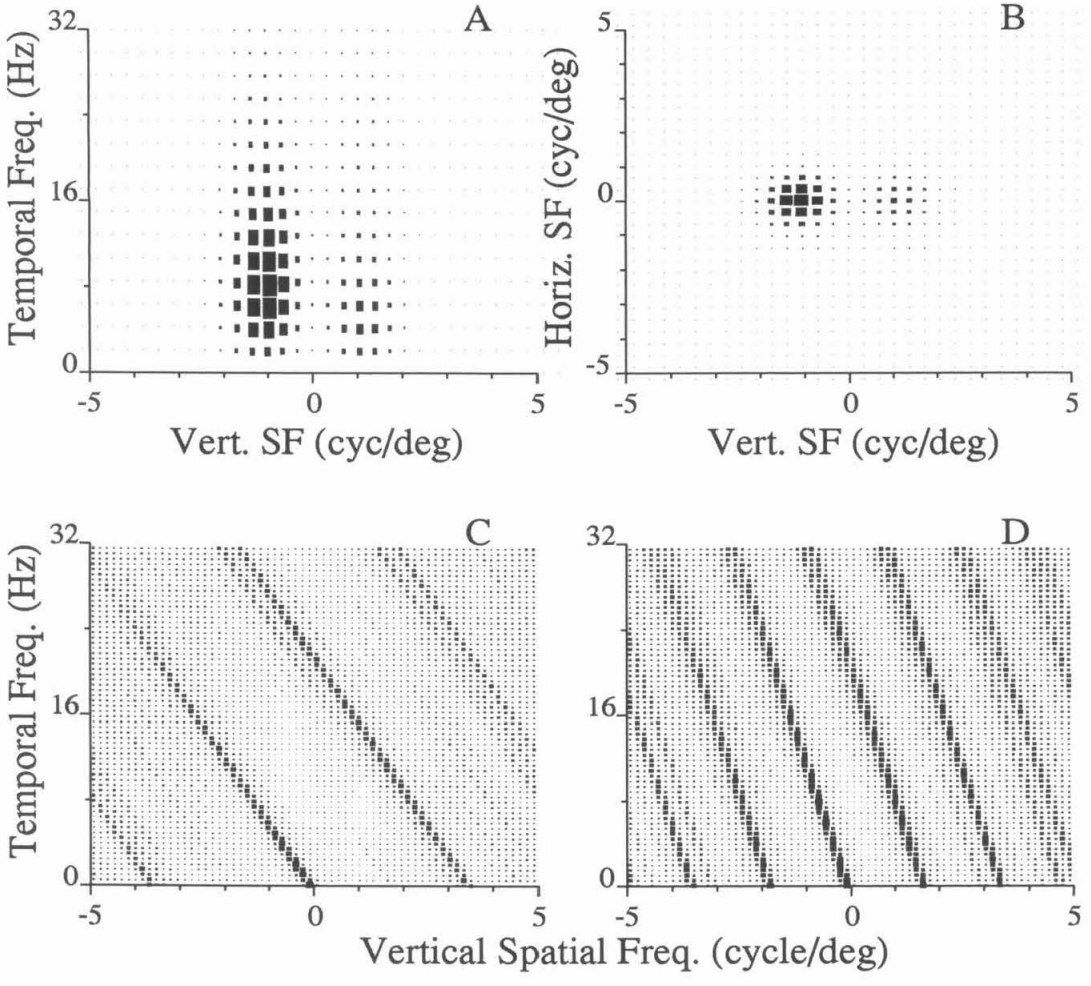

Figure 6.14: Motion energy in the frequency domain. A and B show the results of summing the power spectra of the 3D filters along the horizontal spatial frequency axis (A) and the temporal frequency axis (B). The spatial filter (B) is symmetrical and has narrower band width, as expected for the Gabor function, while the temporal filter has broader bandwidth and is asymmetrical. C and D show the power spectra of $c = 1$ dot stimuli at 5 and 10°/sec, respectively. Copies of the spectrum are replicated because the stimulus is pulsatile. The 45 msec Δt used for the stimulus constrains the first replicated copy of the spectrum to pass through (0,22.2) in C, while the velocity determines the slope. From this, one can see how the spectrum sweeps through the frequency domain, intercepting the filter profiles, as the velocity is changed.

cutoff frequency in the output.

Kelly (1979) measured the human contrast sensitivity to drifting sinusoidal gratings of various spatial and temporal frequencies. The maximum sensitivity at each velocity is redrawn from Kelly's Fig. 15 as a thick line ("max") in Fig. 6.15 here. The peaks of the bands in the power spectrum for our dynamic dot stimulus traveling upward at 10°/sec are shown by dotted lines. For a similar but continuously moving stimulus, no power would fall in the quadrant shown here; this quadrant would have energy from continuous motion stimuli for downward motion only. The striking feature demonstrated by the figure is that there are roughly similar amounts of motion energy in all quadrants even at low velocities and for low spatial and temporal frequencies. The region of the quadrant shown in Fig. 6.15 includes the optimal spatial and temporal frequencies for nearly all V1 cells (Foster et al., 1985) and is a sub-region of Watson and Ahumada's (1983) "window of visibility." Using a two-alternative forced choice paradigm, Watson and Ahumada verified that, at 10°/sec, humans could distinguish sampled from continuous motion at flicker rates as high as 75–150 Hz (see the critical sampling frequency curves in their Fig 5).

Because the dynamic dot stimulus with 45 msec Δt spreads power in oriented bands throughout the region of sensitivity for human psychophysics and macaque V1 neurons, and because the spreading, or filling in, becomes locally more dense at higher velocities, it is possible that the optimal velocities determined for MT cells using this stimulus might be underestimated. Fig. 6.16 shows the preferred speeds for the dynamic dot stimulus for of some of the MT neurons.[1]. Other studies that used smooth motion have reported higher preferred stimulus speeds in MT cells. For example, Zeki (1974) reported that most cells had speed preferences between 5–50°/sec. In the anesthetized monkey, Maunsell and Van Essen (1983) found optimal

[1]The cells used here are those which were recorded with identically seeded stimuli.

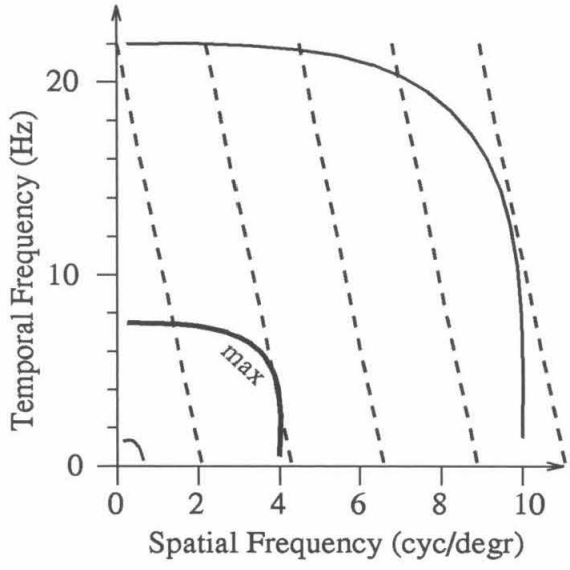

Figure 6.15: The human spatio-temporal frequency threshold surface (solid lines, redrawn from Kelly, 1979) is intersected by the peaks of the power spectrum (dotted lines) of the dynamic dot stimulus. A vertical intercept at 22.2 Hz (see Eqn. 6.24) results from the 45 msec Δt and the slope of the dotted lines is the negative speed. This is a null direction quadrant for the corresponding continuous motion. The peak of the spectrum for continuous motion would be parallel to the lines shown here and would intersect the origin, traversing quadrants II and IV. The thick line shows the maximum human contrast sensitivity for moving gratings at each velocity, i.e., along radii through the origin, while the thin lines show a contour curve approximately one order of magnitude below maximal sensitivity. The entire region here covers most of the optimal sensitivities of V1 cells in macaque monkeys (Foster et al., 1985).

speeds in the range 2–256°/sec (mean 32°/sec). Mikami et al. (1986)[2] found a similar distribution, with a small shift toward faster speeds, in the awake monkey. Rodman and Albright (1987) reported optimal speeds between 5–150°/sec. Lagae et al. (1993) report median velocities near 60°/sec for light and dark bars in MT cells within 23° of the fovea, and they claimed that optimal speed did not depend on eccentricity.

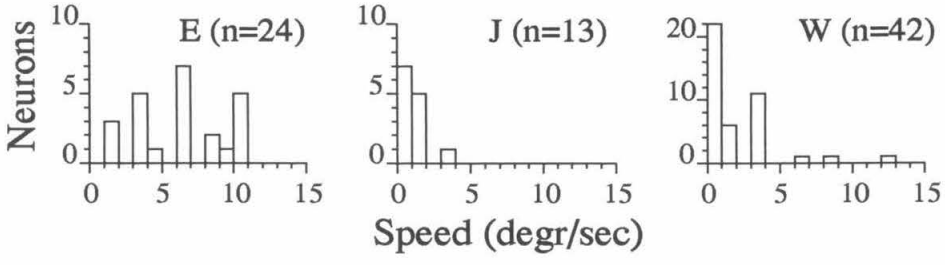

Figure 6.16: Preferred speeds for MT neurons from 3 monkeys (E, J, and W). Most cells preferred speeds below 10°/sec. These low values may be a result of the dynamic dot stimulus having power distributed across the frequency domain, as given by Eqn. 6.21.

Area MT may derive a substantial amount of its input from directionally selective complex cells, and many of these cells are reasonably well modeled by motion energy units (Emerson et al., 1992). Motion energy units are tuned to a particular region of the spatio-temporal frequency domain, as demonstrated with Fig. 6.4. Such units will be excited by the bands of energy from the dynamic dot stimulus; therefore, it is quite possible that V1 cells will be excited that contribute to MT cells at all velocity preferences. Given the frequency domain analysis of the random dot stimulus and the mismatch between the distribution of speeds reported in the literature and speeds used to collect the MT data studied here, there is a reasonable possibility that many of the MT neurons studied here were stimulated with non-optimal speeds for continuous motion. Mikami et al. (1986) state, "... we have observed a suppressive mechanism,

[2]They used stroboscopic motion but changed the Δt to find the optimal speed.

possibly inhibitory, which acts in the *preferred* direction to limit the response of MT neurons to motion of nonoptimal speeds" [original emphasis]. Thus, it is not unreasonable to consider that inhibitory connections onto preferred direction cells may be activated during the coherent $c = 1$ motion studied here. Whether these connections are similar to those postulated to be responsible for the gaps at $c = 1$ reported in Chapter 5 is uncertain.

However, we point out that among the class of dynamic dot stimuli used in the original experiments, the parameters were optimized, and because MT cells are known to have high contrast sensitivity (Sclar, et al., 1990) and show firing rate saturation with only a few dots in their receptive fields (Snowden et al., 1991), it is likely that the firing rates recorded for the $c = 1$ stimuli were as high as those for more globally optimal stimulation, for example, continuous motion—if that is indeed more optimal.

We believe the observation of the banded structure of the power spectrum is important because it may offer a clearer explanation for the stimulus-locked modulation and for some of the psychophysical observations. If the signal dots provide a suboptimal stimulus for a particular neuron, when a more optimal stimulus occurs at random due to noise dots (or signal-noise dot interactions), the neuron may give a strong and transient response. This is consistent with the suggestion from Chapter 3 that the noise dots play a strong role in modulations of the psychophysical response. It is also highly consistent with data shown in Fig. 4.1 in which many of the largest peaks persist for preferred and null direction stimulation, in which signal dots move oppositely but noise dots remain the same. For $c = 1$ motion, when there is very little competition from potentially more optimal stimuli, i.e., there are no noise dots, the firing rate may be determined mainly by the relative strength of activation of the preferred and null direction pathways provided by the banded spectrum of the dynamic dot stimulus. Increasing or decreasing the number of signal dots could change only the absolute level of activation of each pathway. Therefore, fluctuations in the *number*

of signal dots may be removed through a normalization procedure which compares the relative strength of activation across directions of motion. This is consistent with the absence of stimulus-locked temporal modulation for the $c = 1$ stimulus that is reported in Chapter 4. Further tests of these predictions can be performed to some extent with the current database.

6.5 Future Use of the Model

The model described in this chapter has led to the conclusion that a standard motion energy model is consistent with the presence of stimulus-locked temporal modulation observed for incoherent motion. An analysis of the dynamic dot stimulus indicates that it may yield the best responses from MT neurons at slower than their optimal velocities. Under some normalization schemes, this may make the random motion energy contributions from the noise dots responsible for the temporal modulation. This hypothesis can be tested further with the model by including normalization across spatial frequency channels and across directions of motion.

We would like to have a model of an MT cell that could predict the exact locations and sizes of the peaks in the PSTH for the response to a particular dot pattern. More recently, the software that controls the stimulus generation and data collection in the Newsome lab has been modified to save the random seed. We have written software to recompute to dot patterns from the seed. The recomputed stimulus pattern was used as input to the M-unit with various parameter settings and the response is shown at the bottom of Fig. 6.17. Whether there is any correlation between the neuronal output and the M-unit output is debatable. If more data can be collected, or if a more controlled stimulus is used, a model that predicts the neuronal responses may be constructed.

Another aspect of the data which we would like to model is the shapes and pa-

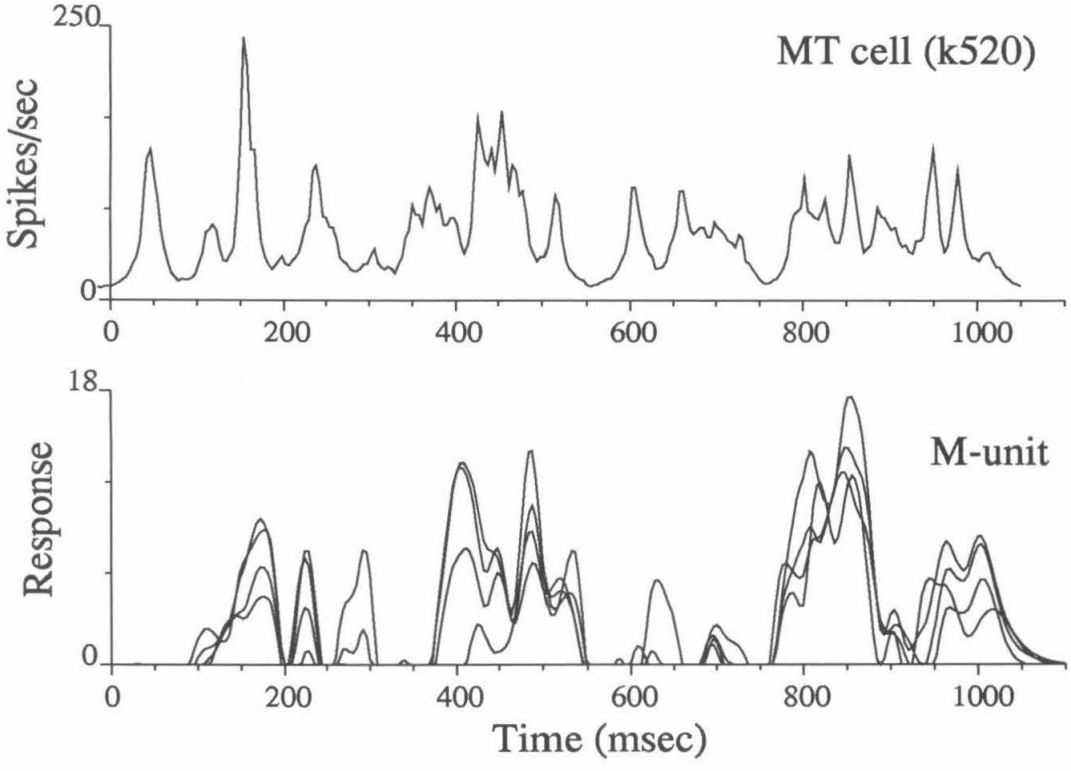

Figure 6.17: The PSTH of neuron k520 in response to a particular $c = 0.13$ stimulus pattern (top) and four superposed responses to the same stimulus from M-units having different parameters (bottom).

rameter dependencies of peaks in the cross-correlations between pairs of MT neurons. Fig. 6.18 shows cross-correlograms that have been computed using the joint post-stimulus time histogram method (Aertsen et al., 1989) for three pairs of MT neurons recorded in a previous study (Zohary et al., 1994). If MT neurons can be described by a collection of motion energy filters, and if neighboring MT cells share similar populations of motion energy sub-units, what predictions does this make about the cross-correlation?

Finally, we would like to model the psychophysical data from the cut experiments in Chapter 3. This will require the model to at least take into account anisotropies of the spatial frequency sensitivity with respect to eccentricity. It is our belief that judicious use of the model can lead to more efficient psychophysical and electrophysiological experimentation.

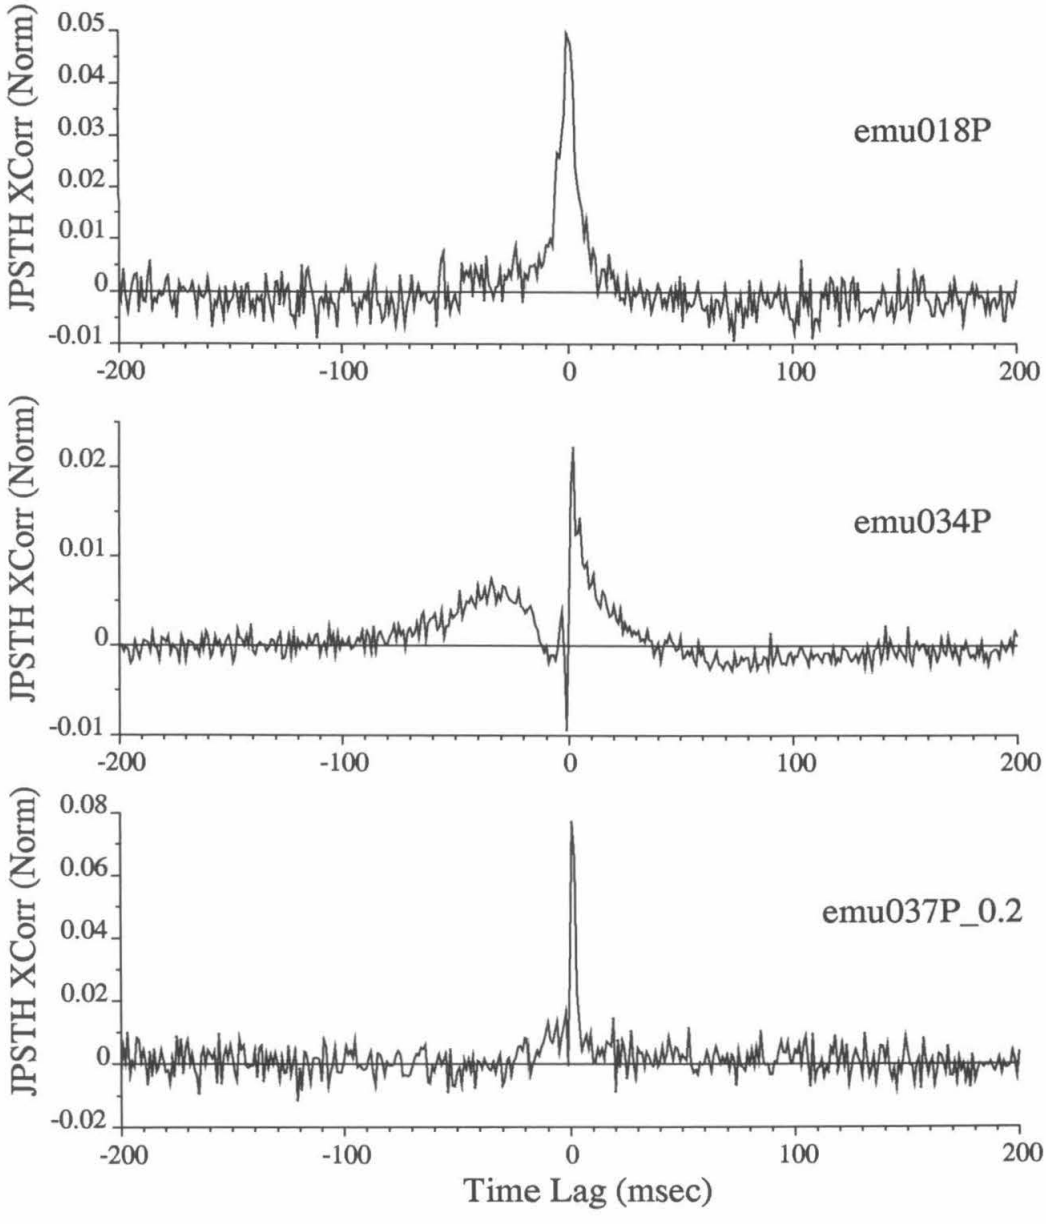

Figure 6.18: Cross-correlations between three pairs of MT neurons. This is the "normalized" version of the joint post-stimulus time histogram method (Aertsen et al., 1989) collapsed (averaged) parallel to the central diagonal. The value is an average correlation coefficient for all bins corresponding to the given lag.

Chapter 7

Power Spectrum Analysis of Bursting Cells

This chapter contains the text from Bair et al. (1994) which summarizes the earliest work reported in this thesis. The introduction and the methods section have been modified, but the rest of the text is presented in nearly its original form. The analytical models of the power spectrum presented here are reworked from a different, more rigorous, perspective in Chapter 8.

7.1 Introduction

It is widely held that visual cortical neurons encode information primarily in their mean firing rates. Some proposals, however, emphasize the information potentially available in the temporal structure of spike trains (Optican and Richmond, 1987; Bialek et al., 1991), in particular with respect to stimulus-related synchronized oscillations in the 30–70 Hz range (Eckhorn et al., 1988; Gray et al., 1989; Kreiter and Singer, 1992) as well as via bursting cells (Cattaneo et al., 1981a; Bonds, 1992). We investigate the temporal fine structure of spike trains recorded in extrastriate area

MT of the trained macaque monkey, a region that plays a major role in processing motion information. The data were recorded while the monkey performed a near-threshold direction discrimination task so that both physiological and psychophysical data could be obtained on the same set of trials (Britten et al., 1992). We identify bursting cells and quantify their properties, in particular in relation to the behavior of the animal.

We compute the power spectrum and the distribution of interspike-intervals (ISI) associated with individual spike trains from 212 cells, averaging these quantities across similar trials. (1) About 33% of the cells have a relatively flat power spectrum with a dip at low temporal frequencies. We analytically derive the power spectrum of a Poisson process with refractory period and show that it matches the observed spectrum of these cells. (2) About 62% of the cells have a peak in the 20–60 Hz frequency band. In about 10% of all cells, this peak is at least twice the height of its base. The presence of such a peak strongly correlates with a tendency of the cell to respond in bursts, i.e. 2–4 spikes within 2–8 msec. For 93% of cells, the shape of the power spectrum did not change dramatically with stimulus conditions. (3) Both the ISI distribution as well as the power spectrum of the vast majority of bursting cells is compatible with the notion that these cells fire Poisson distributed bursts, with a burst-related refractory period. Thus, for our stimulus conditions, no explicitly oscillating neuronal process is required to yield a peak in the power spectrum. (4) We found no statistically significant relationship between the peak in the power spectrum and psychophysical measures of the monkeys' performance on the direction discrimination task. (5) For cells firing bursts, ROC analysis shows that the "event" rate, where an event is either a single burst of spikes or an isolated spike, is on average a more sensitive measure of visual stimulus direction than the total number of spikes, used previously (Britten et al., 1992), implying that the number of spikes in a burst is less stimulus-dependent than the overall firing rate or the rate of bursts.

Only a few papers have focussed on the possible significance of bursting for neuronal coding. Cattaneo et al., (1981a,b) report that complex (but not simple) cells in area 17 of the anesthetized and alert cats frequently respond with bursts. They show that the frequency of bursts (or "grouped spikes") varies strongly with the spatial frequency and orientation of sinusoidal drifting gratings, while the frequency of "isolated spikes" only weakly depends on these parameters, encoding rather the contrast of the stimulus. Bonds (1992) found in his analysis of cat striate neurons that the structure of spike trains—in his case the frequency and duration of bursts—can vary substantially on the basis of how the firing rate was generated. For instance, presentation of stimuli at non-optimal orientations at high contrasts yields bursts that are shorter than those generated by lower contrast stimuli at optimal orientations. Legéndy and Salcman (1985) hypothesized functional significance for burst firing patterns in spontaneously active striate neurons of alert cats, but by their definition, bursts included mostly long periods (0.5–2.0 sec) of significant elevation in firing rate. Finally, Crick (1984) postulated that the neuronal expression of selective visual attention is the production of bursting in a subset of thalamic neurons. This bursting, in combination with a short-term and transient alteration in the synaptic strength (as proposed by von der Malsburg, 1981), could lead to the short-term formation of transient cell assemblies at the level of cortex (see also Crick and Koch, 1990, 1992).

We now analyze this single-cell database with an eye towards describing the temporal structure of MT responses and uncovering any relationship between the temporal structure and the psychophysical performance of the animal. We only consider data from well-isolated single neurons. We find that the temporal structure of MT responses is characteristic for each neuron and does not change in a stimulus-specific manner. Spikes are distributed almost randomly in time for some neurons, but are highly non-random for other neurons, being characterized by occasional "bursts" in which spikes are tightly clustered in time. Both firing patterns are well-described

by simple models that require no intrinsically oscillatory process. A signal detection analysis indicates that bursting neurons convey more information about the direction of motion in the stimulus if bursts are considered to be individual signaling events.

7.2 Data Analysis

From the initial database of 216 cells, four are not considered here. Three of these were not recorded for the full 2 sec (*w044*, *w045*, *w046*), and the fourth cell showed an abrupt change in firing rate during the experiment and was presumed to be damaged (*j036*).

The experimental methods are described in Chapter 2, and only the computational methods are described here. Given the importance of applying the underlying mathematical transformations correctly, we justify our analysis in detail.

7.2.1 Computation of Power Spectra

For a real and continuous function $F(t)$ observed between $t = -T$ and $t = +T$, the associated continuous Fourier transform $\tilde{F}$ at the frequency f is given by

$$\tilde{F}(f) = \int_{-\infty}^{+\infty} e^{-i2\pi ft} F_T(t) dt, \tag{7.1}$$

where $F_T(t) = F(t)$ for $|t| \leq T$ and 0 outside the observation interval. The autocorrelation function associated with F_T is given by

$$R(t) = \lim_{T \to \infty} \frac{1}{2T} \int_{-T}^{+T} F_T(t_1) F_T(t_1 + t) dt_1. \tag{7.2}$$

Following the Wiener-Khintchine theorem, the Fourier transform of the autocorrelation is equal to the power spectrum, that is

$$S(f) = \int_{-\infty}^{+\infty} e^{-i2\pi ft} R_T(t) dt, \tag{7.3}$$

or,

$$S(f) = \lim_{T \to \infty} \frac{1}{2T} \tilde{F}(f) \tilde{F}^*(f), \tag{7.4}$$

where * denotes the complex conjugate. It can be seen that the power spectrum is always real and symmetric.

However, given the discrete sampled nature of our spike trains, we require the use of the discrete Fourier transform with the two associated problems of **(i)** the variance inherent in the estimate of the discrete power spectrum and **(ii)** aliasing due to a finite sampling interval. We perform Fourier transforms on spike trains using the standard Fast Fourier Transform (FFT) algorithm and compute one-sided estimates of the power spectral density using overlapping data segments and windowing (Press et al., 1988). To emphasize the difference between the true power spectrum $S(f)$ and the one we compute on the basis of the sampled data, we denote the latter by $S'(f)$. The data submitted to the FFT algorithm is a sequence of 1's and 0's, where each 1 represents an action potential in a spike train sampled at 1 kHz. Transforming a 2 sec long spike train yields a one-sided spectrum with a frequency resolution of about 0.5 Hz from 0 up to the Nyquist frequency of 500 Hz.

Because we are only interested in studying broad trends in the data over a relatively wide band of frequencies, we do not require such high frequency resolution. Furthermore, the variance associated with the estimation of the power spectrum can be reduced by using larger frequency bins, i.e. by sacrificing frequency resolution. Thus, we break the 2 sec long trial into smaller segments, typically using 12 overlapping data segments of 256 msec duration, thereby utilizing 1664 msec of the 2000 msec

spike train. We always begin the first segment at 336 msec to eliminate from analysis the transient response to the onset of the stimulus; however, we find essentially no difference in the results when the analysis is performed using the entire 2 sec spike train (not reported here). For each 256 msec segment, a two-sided power spectrum was computed. Given the fact that the power spectrum is always symmetric, we normalize the two-sided spectrum to a one-sided spectrum with 128 entries lying at equally spaced intervals between 0 and 500 Hz (with $\Delta f \approx 4$ Hz).

To further improve the spectral estimate, a triangular Parzen window was applied to each segment to reduce spectral leakage arising due to the finite duration of spike trains (Harris, 1978). Since windowing would otherwise throw away data at the ends of each segment, it becomes important to use overlapping segments. In addition to the averaging due to data segmentation and symmetry for a single 2 sec long trial, we usually—except if otherwise stated—average the power spectra over all trails for a particular c for an individual cell. Since we average in the frequency domain, our method is not sensitive to the exact phase relationship of the response with respect to stimulus onset. When computer generated data is shown, averages are over the equivalent of 1000 trials of 2000 msec duration each.

While we can reduce the variance in the spectra by averaging, we cannot avoid aliasing due to temporally sampled data. To what extent does aliasing play a dominant role in shaping our spectra? A continuous abstraction of a spike train is a set of occurrence times for action potentials which are idealized as Dirac $\delta(t)$ impulse functions. In recording these occurrence times, the continuous function is not sampled in the usual manner, rather the action potentials are shifted to nearby sampling points. If this were not the case, then most action potentials would be missed altogether. This type of sampling is a form of data binning in which each bin has width equal to the sampling interval and the assumption is made that at most one event occurs per bin. Given a sampling interval of 1 msec, this is a reasonable assumption for

most neurons. We present a modification of Schild and Schultens' (1986) analysis of aliasing in binned data for post-stimulus time histograms to understand what effect this has on the spectrum.

Figure 7.1 summarizes this analysis for a special case which has a simple and revealing solution. The left side illustrates how the binning process that converts the continuous spike train (in **a**) into sampled data (in **e**) is described as a convolution with a boxcar function (in **b**) followed by multiplication with a comb function (in **d**). We assume the special case where the continuous spike train is totally random (Poisson) and therefore has a flat power spectrum (in **A**) with a delta function at the origin (which we ignore here; see Modeling section). Convolving with the boxcar in the time domain corresponds to multiplying the power spectrum by the square of a *sinc* function ($sinc(x) = \sin(x)/x$; in **C**). Multiplying by the comb function in time corresponds to convolving the power spectrum with a comb with inverse spacing, which replicates and sums copies of the $sinc^2$ function at 1 kHz intervals (Fig. 7.1**C**). The sum of the original $sinc^2$ plus the infinite number of $sinc^2$ functions shifted by 1 kHz is a flat spectrum, which is what we started with; however, only the copy of the $sinc^2$ centered at zero is contributing to the true power spectrum while the other shifted copies are contributing aliased frequencies (Fig. 7.1**E**). In this case, frequencies below about 200 Hz are relatively uncorrupted and this is the frequency range which concerns us. As discussed later, the power spectrum of most cells in the database is relatively flat, especially at higher frequencies. The worst cases of aliasing occur for neurons with interspike intervals at least an order of magnitude more regular than we observe here (Schild and Schultens, 1986).

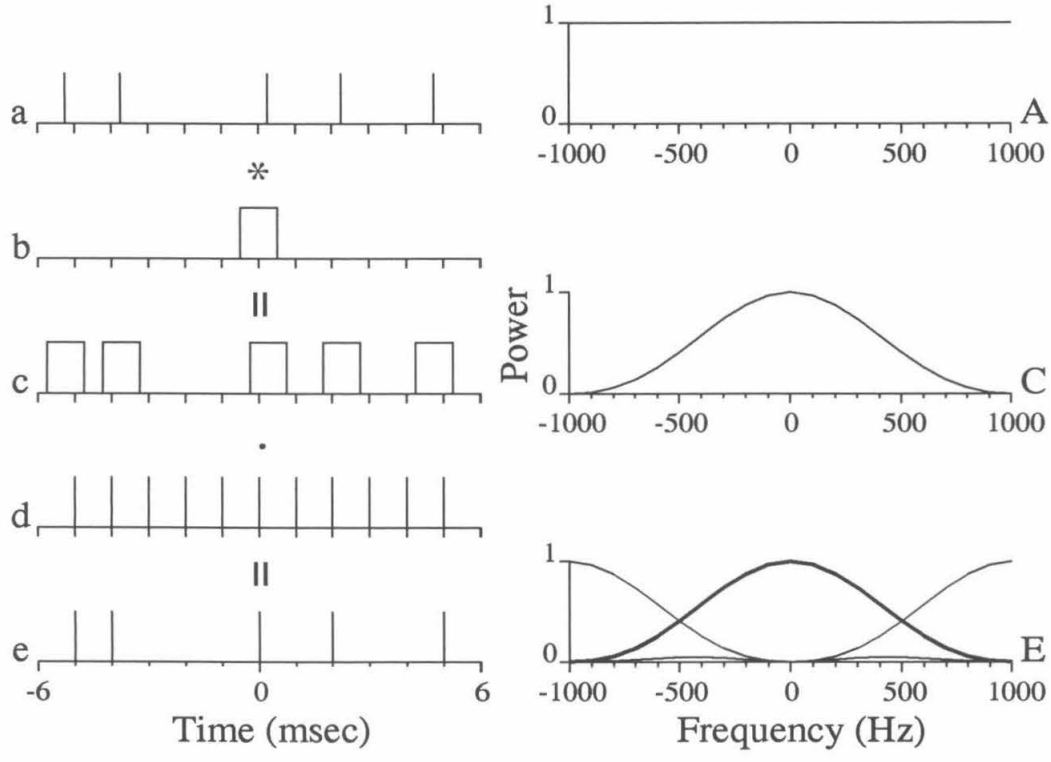

Figure 7.1: Aliasing in the power spectrum $S(f)$ of a spike train. The ideal spike train (a) is a sequence of δ-functions at arbitrary time instants, but in recorded spike train data (e), spikes are assigned to sampling points. To go from (a) to (e) we convolve with a 1 msec wide binning function (b) to yield (c) and then multiply by a 1 kHz sampling function (d) to yield (e). On the right side are power spectra for the special case where (a) is random (Poisson). Ignoring the DC component, (A) is flat. Convolving with (b) multiplies the power spectrum by the square of the *sinc* function, yielding (C) as the spectrum of (c). Sampling in this special case corresponds to adding up infinitely many copies of the $sinc^2$ at 1 kHz intervals (E) which yields a flat power spectrum (not shown) like (A). The dark curve in (E) demonstrates the contribution to the final power spectrum from the true spectrum, while the light curves (shifted $sinc^2$ functions) show the contribution from aliasing. There is relatively little aliasing for $|f| < 200$ Hz.

7.2.2 Quantification of Spectrum Shape

In the Results section, we utilize a scheme for classifying cells based on the shape of their power spectra (for a related approach, see Ghose and Freeman, 1992). The classification of a cell depends upon the classification of its set of power spectra, one at each c value. Here we describe in detail how to compute the set of spectra for a cell, and how to classify the cell based on that set.

The estimated power spectral density for a cell at a given c, $S'_c(f)$, is computed as the average of the spectra for all spike trains recorded at that c value normalized by their average spike rate. For this normalization, the computation of average spike rate must take into account the effect of the overlapping data windows and the shape of the Parzen window. These corrections are particularly important for spike trains with very few spikes. If the individual spike trains were Poisson, then the expected value of each normalized spectrum would be 1.0 at all non-zero frequencies. This normalization reduces the variance in S'_c due to fluctuations in spike rate from trial to trial while preserving relationships between peaks and dips within each spectrum.

As discussed in the Results section, many cells have peaks in their spectra in the region below 60 Hz (see Fig. 7.2a and b, bottom row), and we wish to measure the height of that peak with respect to the associated dip or lowest baseline level at higher frequencies (see arrows on $S'(f)$, Fig. 7.2 cell **a**) and at the same time determine if the peak is present for all c. The location of the peak (if one exists) is computed by sliding a fixed-width window along the set of spectra while looking for the frequency f_p (at the center of the window) which maximizes the integral within the window summed over all S'_c. The constraint 20 Hz $< f_p <$ 60 Hz is used to keep the window away from the peak at $f = 0$ (i.e. the DC component) and to avoid scanning higher frequencies where peaks are absent. The average value of S'_c within the window will be called the peak level, P_p. A second sliding window is used in a

similar, but minimizing, procedure to find the lowest point, or baseline level, at higher frequencies. The location of the lowest point is quantified by f_b, the center of the window, and the average value within the window, P_b, will be called the baseline level. We constrain f_b to be greater than f_p and less than the cutoff frequency, 500 Hz. All windows are seven spectral bins wide (3.9 Hz/bin) and therefore sacrifice accuracy of localization for noise immunity.

Once the peak frequency f_p and the baseline frequency f_b are determined, a cell is classified as a *burst* cell (later we explain the link between peaks and burst firing) if for at least 90% of all c values $P_p > P_b$, i.e., if at nearly all coherence levels, the spectrum has a peak in the 20–60 Hz range. If a cell fails to be classified as a burst cell, then we attempt to localize a dip in the power spectrum (see arrow on $S'(f)$, Fig. 7.2 cell **d**), bottom left, which has a dip near 20 Hz) using another sliding window. We take f_d to be the center frequency of the seven bin window which minimizes the integral within the window summed over all c, where 20 Hz $< f_d <$ 500 Hz. We take the average value of the spectra within a window centered at f_d to be P_d. We classify a cell as *nonburst* if for at least 90% of all c values $P_d < 1.0$, i.e., if at nearly all coherence levels the spectrum has a dip below the expected baseline level (which manifests itself at high frequencies) for a Poisson-like spike train, which is 1.0 due to our spike rate normalization. This definition would result in classifying a cell with Poisson distributed spikes (that therefore has a flat power spectrum) as neither burst nor nonburst, but since all cells studied here show evidence of refractory periods, this case does not occur in practice. Note that for a pacemaker cell "oscillating" in the 20–60 Hz band, P can become arbitrarily large as the oscillation becomes increasingly regular.

If a cell fails to be classified as either burst or nonburst it is classified as *mixed* since at some c values it lacks a significant peak, while at others it lacks a significant dip, in the 20–60 Hz range. To avoid classifying a cell based on too little data, trials

with less than 5 spikes in the 336–2000 msec time window are discarded, and c values with less than 8 valid trials are not represented in the set of spectra. Finally, a classification is made only when there are at least 3 different c values with valid S'_c. Typically, there are six S'_c for $c < 0$, six for $c > 0$ and one for $c = 0$.

A measure of the shape of the power spectrum, P, is associated with each classified neuron. For burst cells, the ratio of the peak to the baseline is used, $P = P_{burst} = P_p/P_b$. For nonburst cells, the ratio of the dip level to the ideal baseline, 1.0, is used, $P = P_{nonburst} = P_d$. The value for mixed cells depends on the subcategorization, i.e. P follows the definition for burst if S'_c had a peak for the majority of c values but follows the definition for nonburst otherwise. When discussing the shape of power spectrum, we will simply refer to P when the particular definition is understood from context.

7.2.3 Other Methods

The post-stimulus time histograms (PSTH) are computed from the single trial data by averaging over all trials with identical stimulus conditions, using a bin width of 10 msec. They are normalized to show spike rate rather than counts per bin. Inter-spike interval (ISI) distributions are computed with 1 msec bin width. Power spectra, $S'(f)$, are usually normalized to match continuous spectra under the assumption that spikes can be described as Dirac δ-functions. In this case, the vertical offset is roughly proportional to the spike rate, and for nonburst cells, the flat level at higher frequencies is usually an accurate reflection of the spike rate, as in Fig. 7.12. Under the second spike-rate normalization (discussed above), spectra are divided by the average spike rate so that all are nearly the same height to allow comparison of shapes, as in Fig. 7.5.

Vertical truncation of histogram plots is indicated by open histogram bars near

the top of the graph.

7.3 Experimental Results

We begin by describing the population of MT cells with respect to two statistical measures of the temporal fine structure of spike trains: the ISI distribution and the power spectrum $S'(f)$. The first measure is an order-independent statistic since it contains no information about the temporal order in which the intervals occur. For instance, all short intervals could have occurred at the beginning of each trial and all long intervals at the end, or each short interval could have been followed by a long one. The estimate $S'(f)$ is order-dependent since it depends on temporal relationships between events at scales beyond single intervals. Although many different $S'(f)$ may be associated with a particular ISI distribution, we find for this database that the shape of the ISI distribution predicts the shape of the power spectrum quite well, and that the tendency of a cell to fire bursts of action potentials is the basis for the prediction.

Estimates of the ISI and $S'(f)$ are shown in Fig. 7.2 for four cells from the database. Segments of typical spike trains from each cell are shown at the top, and below them, the PSTHs show that the average firing rate is relatively constant throughout the period over which we compute the ISIs and power spectra, from 336–2000 msec. Although not shown here, including the initial transients had little effect on the shape of $S'(f)$. The autocorrelation functions, $R'(t)$, (not shown) for these cells do not show ringing, even when the associated spectrum $S'(f)$ has a prominent peak of the type seen in Fig 7.2 (cell **a**, bottom). Only a single cell showed strong ringing in $R'(t)$ and this was for $c = 1$.

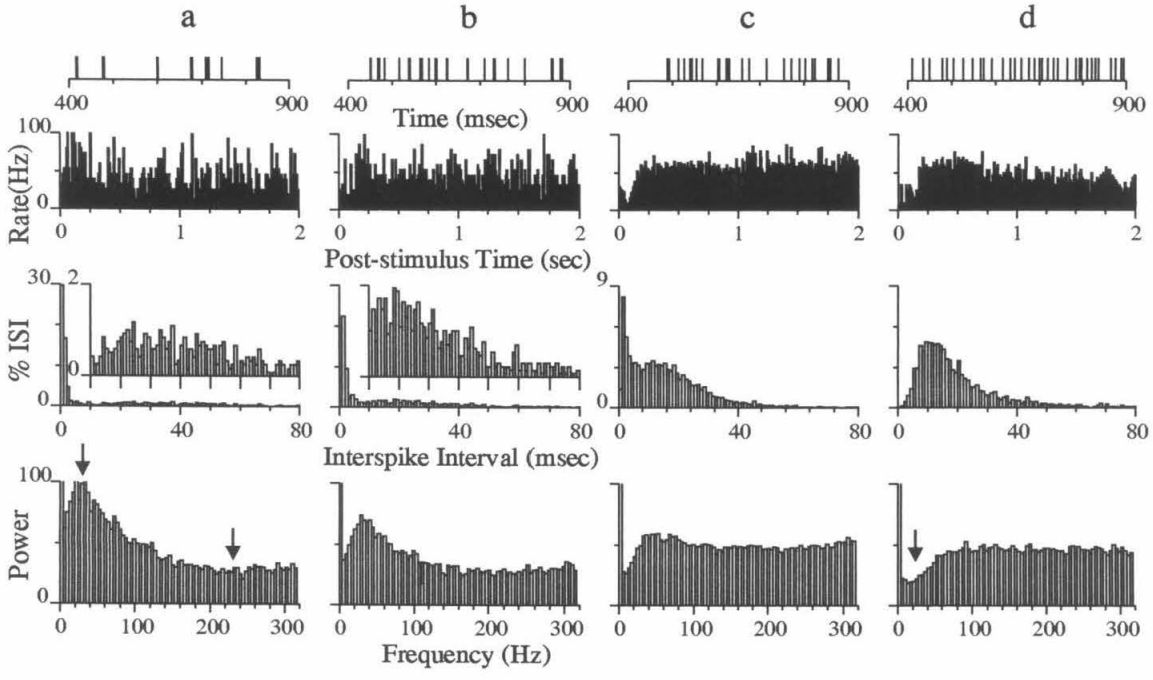

Figure 7.2: Spike train statistics for four area MT neurons responding to $c = 0.256$ motion. Each column corresponds to a different cell, and the cells are arranged from most bursting (a) to least bursting (d). The top row shows 500 msec segments of spike occurrence times. The distribution of intervals between spikes are shown by the ISI histograms in third row. Over 50% of intervals fall in the 1,2 or 3 msec bins in ISI (a) while less than 2% do so in ISI (d). (Note scale change for ISI (c) and (d).) For (a) and (b) ISI insets (horizontal axes begin at 10 msec and remain aligned with main histogram) expand the vertical axis to show the peak near 20–30 msec which corresponds to intervals between bursts. Power spectral densities, $S'(f)$, (bottom row) have peaks in the 20–60 Hz range for the cells (a,b,c) which have peaks in the ISI below 4 msec. The dip in $S'(f)$ at low frequencies (arrow in (d)) indicates the presence of a refractory period for spikes or bursts. All four cells fire at roughly similar average rates (40–60 spikes/sec, see PST histograms, second row) and the response is relatively maintained following the transients during the first 200 msec of stimulation. Plots are based on averages from 15 trials (a,b), 60 trials (c), and 30 trials (d). Frequent burst firing causes PST (a) and (b) to be excessively noisy.

7.3.1 Bursting Cells

We were intrigued by the persistent tendency of certain cells to burst frequently, that is to discharge a group of 2–4 tightly clustered action potentials with interspike intervals of no more than 3 msec (see, in particular, the ISI histogram in Fig. 7.2a and b). Cells that respond frequently in this manner to the visual stimulus show a bimodal ISI histogram with a pronounced peak at short, i.e. 1–3 msec, intervals and a second, much lower and broader, peak at longer intervals (see ISI insets in Fig 7.2, cells **a** and **b**). The first peak is caused by the interval distribution within a burst, while the second peak is partly due to intervals between bursts (interburst intervals) and partly due to intervals between isolated spikes.

Figure 7.2 illustrates the trend in burstiness observed in our data, from strongly bursting (cell **a**), to weakly bursting (cell **c**) and non-bursting cells (cell **d**). Figure 7.2 also reveals a second and correlated trend, that some cells have a peak in their power spectrum between 20 and 60 Hz (cells **a** and **b**, somewhat in cell **c**), while others (cell **d**) have a spectrum that is flat with a dip at low frequencies (except at the origin $f = 0$). The mean center frequency of the peak in the estimated spectrum S' for all cells with such a peak is 41 Hz (with a standard deviation of 9 Hz), so in the remainder of this manuscript, we will refer to this frequency range as the 40 Hz band, with the understanding that the location of the peak varies from cell to cell. Of 212 cells analyzed, 71 had relatively flat averaged spectra, $S'(f)$, with a dip at low frequencies, and 131 had peaks in the 40 Hz range of their averaged spectra. The remaining 10 cells had too little data to judge accurately the shape of the spectrum.

As described in the Methods section, we classified the cells as either **burst** or **nonburst** based on the shape of their power spectra using a scheme that takes into account the possibility that the spectra might change as the stimulus motion coherence c varies. The criteria for this classification were designed to be strict so that

cells which did not always, regardless of stimulus direction and coherence, behave as burst or nonburst would be classified as **mixed**. Of 212 cells, 10 were eliminated from classification because they did not meet the minimum standards for number of spikes, trials, and coherence levels. Of the remaining 202 cells, 125 were classified as burst, 61 as nonburst, and 13 as mixed (in 3 cases, visual inspection disagreed with the classification algorithm, so these cells were dropped from consideration). It is striking that 93% of the cells were classified as either burst or nonburst because it indicates that this rough categorization is stimulus invariant.

As also discussed in the Methods section, we define a measure P of the shape of the power spectrum. For burst cells, P is defined as the ratio of the height of the peak in the 40 Hz range to the level of the baseline dip at higher frequencies (see arrows on $S'(f)$, Fig. 7.2 cell **a**). For nonburst cells, P is defined as the ratio of the level of the dip at low frequencies (see arrow on $S'(f)$, cell **d**) to the flat level at higher frequencies. For an ideal Poisson cell, $p = 1$ regardless of the classification as burst or nonburst; for a pacemaker cell "oscillating" in the 20–60 Hz band, P becomes arbitrarily large as the oscillation becomes increasingly regular.

With P as a measure of the shape of S', we developed a measure of burstiness based on the ISI. Given the well known distinction between bursting and non-bursting cells based on intracellular current injections in rodent slice preparations (McCormick et al., 1985; Connors and Gutnick, 1990), we attempted to find a metric that would classify all of our cells into two (or more) segregated groups according to the degree of burstiness. For this purpose, we introduce the measure B as the percent of the ISI histogram in the 1, 2, and 3 msec bins.

The variable B is similar to other proposed measures of burstiness which are based on the proportion of the ISI distribution below a cutoff value (Cattaneo et al., 1981b; Abeles, 1982). We also considered another measure of burstiness, $\tilde{B}$, based on the ratio of the number of intervals in the 2 msec ISI bin to the 5 msec ISI bin. This

variable has the potential advantage that it is able to distinguish between a bursting cell with a bimodal ISI histogram and a very fast firing cell that has a unimodal ISI histogram concentrated below about 10 msec. However, $\tilde{B}$ is quite sensitive to fluctuations in the trough between the peaks of a bimodal histogram, and its value is less stable. We will use B as the measure of burstiness here but point out that B and $\tilde{B}$ tend to be highly correlated, at least for our database.

Figure 7.3 shows the frequency distribution for these three statistical measures for all cells averaged over all stimulus conditions. We interpret the histogram for P to represent a unimodal distribution. The dip at unity is an artifact of our classification system because P is based on regions of the power spectrum that are chosen for maximizing the peak-to-trough ratio or minimizing the trough-to-baseline ratio for burst and nonburst cells, respectively. The long left tail of the distribution for B shows that many cells have less than 1% of their intervals shorter than or equal to 3 msec, such as cell **d** of Fig. 7.2. The distribution for $\tilde{B}$ is spread over many orders of magnitude and shows a hint of bimodality. Overall, however, it is difficult to segment the data into two classes based on these histograms, since many burst and nonburst cells fall in overlapping regions in the histograms for B and $\tilde{B}$. We stress, therefore, that the burst and nonburst classifications are primarily tools for defining two ends of what appears to be a continuum.

For burst cells, P changes relatively little with stimulus condition and appears to reflect primarily an intrinsic property of these cells in an alert and trained monkey. As we show next, in such cells P is highly correlated with B. For nonburst cells, B often changes systematically with spike rate and is therefore not as revealing about intrinsic properties.

The close connection between bursting and the shape of the power spectrum is illustrated in Fig. 7.4. Here the value of B for individual cells is plotted against the associated P. The values shown here are averaged over trials at $c = 0$ and trials at

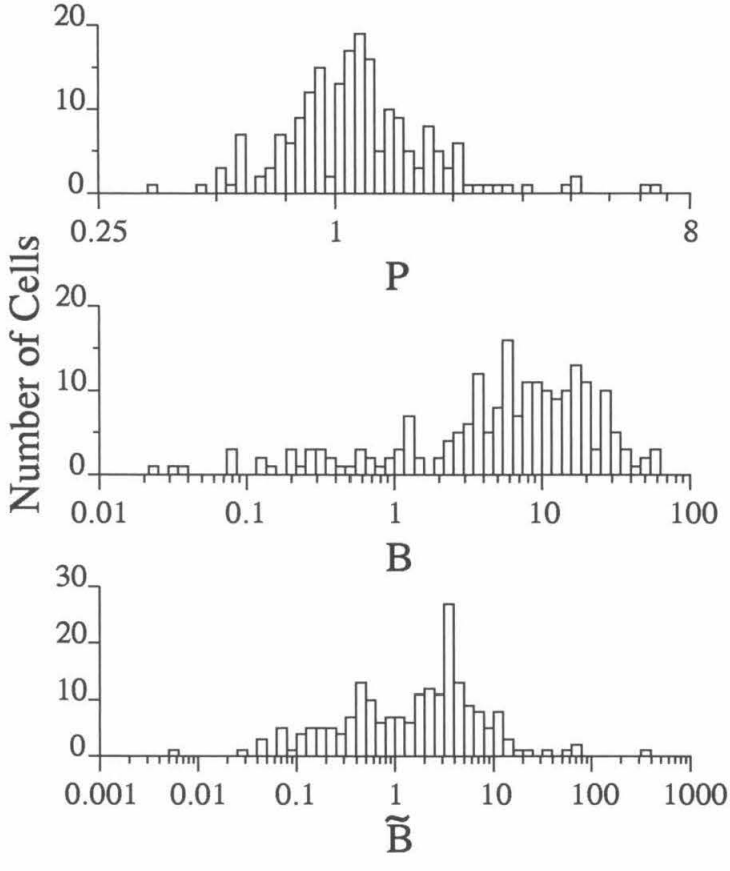

Figure 7.3: Frequency histograms for our database of 212 MT cells for P, the measure of the shape of $S'(f)$, and two measures of burstiness, B and $\tilde{B}$, averaged over trials for all c. The upper two plots show that the distribution of cells with respect to the statistics P and B (the fraction of the ISI in the 1, 2 and 3 msec bins) is primarily a continuum. The dip near 1.0 in the distribution of P is an artifact of the classification of cells as burst or nonburst (see text). All burst cells have $P > 1.0$, nonburst cells $P < 1.0$, and mixed cells that are neither one or the other have P near 1.0. The distribution of B has a long left tail due to cells which rarely fire a second spikes within 3 msec. The bottom plot shows the distribution for a second measure of burstiness, $\tilde{B}$, the ratio of the 2 msec bin to the 5 msec bin in the ISI histogram. B and $\tilde{B}$ are highly correlated for our database, but we use B since it is less sensitive to noise than the ratio measure.

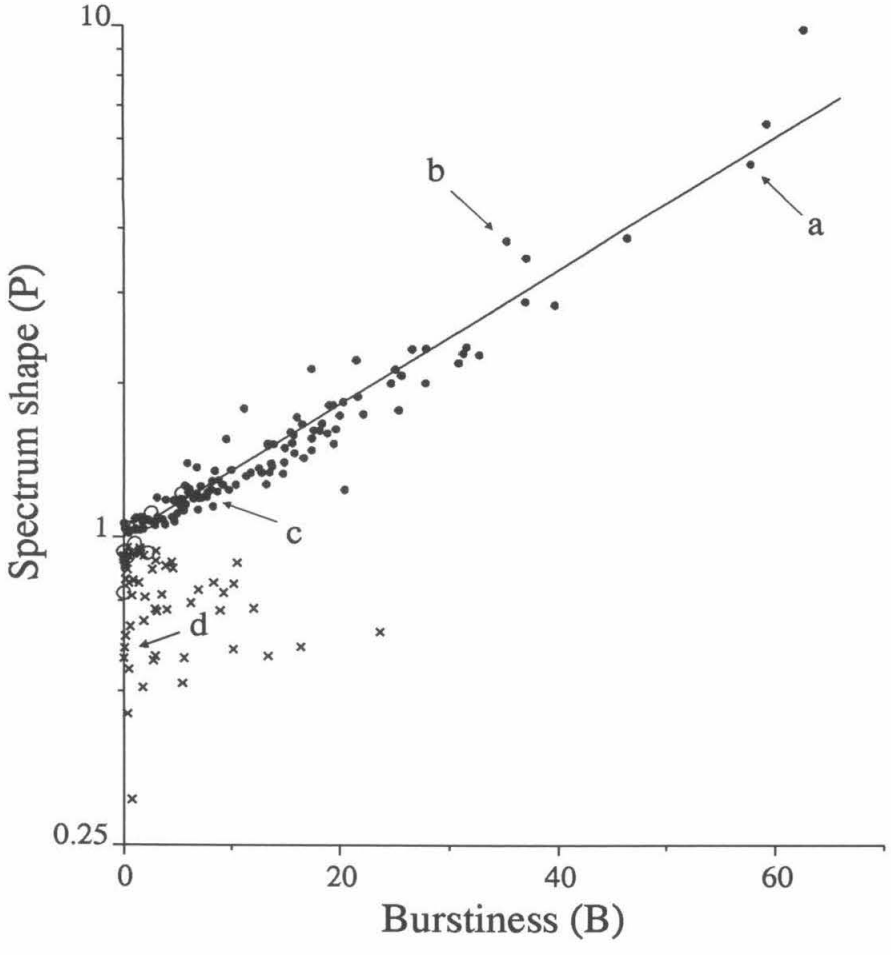

Figure 7.4: $\log(P)$ is plotted against B for burst (filled circles), nonburst (x's), and mixed (large open circles) cells. Among burst cells, the correlation is strong between $\log(P)$ and B, as shown by the fit from linear regression (line). Hypothetical cells firing Poisson spike trains would fall along the horizontal line $P = 1$, separating the burst and nonburst cells. Mixed cells straddle this line and are neither strongly burst nor nonburst. Since B often has a small but systematic variation with spike rate, values shown here are computed by averaging only over trials with responses that are statistically indistinguishable from those at $c = 0$. Results are very similar when values are averaged over trials for all c. The arrows show the points corresponding to the four cells of Fig. 7.2.

other low c values for which the monkey's and the neuron's responses were statistically indistinguishable from responses at $c = 0$. The three groups, burst, nonburst, and mixed, are plotted together in Fig. 7.4. For the burst cells (filled circles), there is an obvious strong correlation between $\log P$ and B, with the line indicating an empirical, exponential fit: $P = e^{0.03B}$. Thus, the more a cell tends to fire action potentials in tight bursts rather than as isolated spikes, the higher the peak in its power spectrum in the 40 Hz band. We will explore the reasons for this behavior further in the Modeling section. Cells with purely Poisson-distributed spikes would fall along the horizontal line $P = 1.0$, with low firing rates near the origin, and higher firing rates to the right. The strong correlation evident here between B and P justifies classifying cells as either "burst" or "nonburst" based on the shape of the power spectrum.

7.3.2 Relation of the 40 Hz Peak to Prior Measures of Neuronal and Psychophysical Performance.

Previous analyses of this data set have identified several interesting parallels between the psychophysical performance of the monkeys and the responses of single MT neurons. In all of these analyses, the response of a neuron was considered to be the total number of spikes occurring during the period of visual stimulation (Newsome et al., 1989a,b; Britten et al., 1992). We now consider to what extent temporal structure, here the presence and amplitude of the 40 Hz peak in the power spectrum, reflects any aspect of the visual stimulus or the monkey's behavior. Specifically, we ask the following questions: 1) Does the prominence of the peak vary with the strength of the motion signal, c? 2) Is the peak affected by the behavioral state of the animal? 3) Does the spectral peak develop or change with time during the course of a two second long trial? 4) How is the peak correlated to prior measurements of cell sensitivity based on total spike counts? 5) Is the size of the peak correlated with the monkey's

decisions concerning direction of motion for a particular stimulus condition?

We suggested in a previous section that the burstiness of a cell, quantified by the ratio P, is largely independent of the strength of the motion signal in the visual stimulus c. Qualitatively, this point is supported by the spectra illustrated in Fig. 7.5, computed for one burst and one nonburst cell, which appear fairly constant in shape for varying values of c, and by Fig. 7.6 which shows the mean and standard deviation for P as a function of c for a burst and nonburst cell.

To analyze the relationship of P and c quantitatively, we first conducted a one-way analysis of variance (ANOVA) for each neuron to determine whether P varied significantly within the range of c's tested. 118 of 202 neurons (58%) failed to show heterogeneity by this test ($p > 0.05$), and we conclude that P is completely independent of c for these cells. For the remaining neurons we performed a multiple regression analysis to determine whether c influenced P in a systematic manner. For the great majority of our neurons, the mean response increased with c, causing changes in the shape of the power spectrum that are related to the presence of the refractory period (see Modeling Section below). We therefore included mean firing rate as a co-regressor in our multiple regression analysis to disentangle the effects on P of c and mean firing rate.

Only 20 of 202 MT neurons (10%) showed a significant relationship of P and c (multiple regression, $p < 0.05$); the slope of this relationship was negative for 13 cells and positive for the remaining 7 cells. For all neurons, furthermore, the slope of the regression line relating P to c was sufficiently small that the classification of a cell as burst (mean $P > 1$) or non-burst ($P < 1$) was unambiguous. It appears, therefore, that c has no strong or systematic impact on P for our population of MT neurons considered as a whole. In most of our subsequent analyses, therefore, P is averaged across c to obtain a single index of burstiness for each cell.

Is the spectral peak influenced by the behavioral state of the animal? To answer

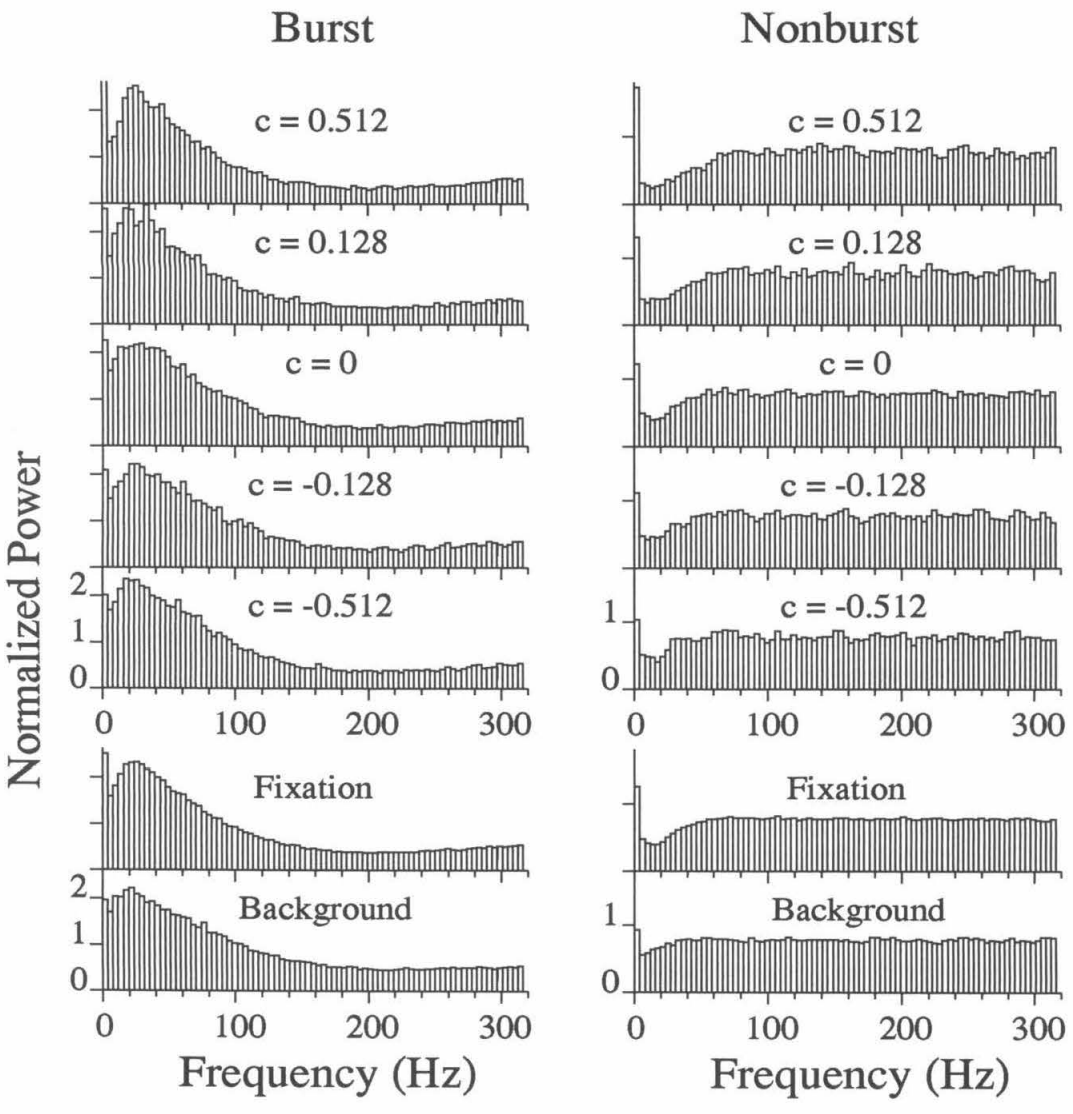

Figure 7.5: Power Spectra, $S'(f)$, are shown for motion coherence values ranging from highly coherent preferred direction motion ($c = 0.512$) to highly coherent null direction motion ($c = -0.512$) for the burst cell (a) and the nonburst cell (d) from Fig. 7.2. Spectra are also shown for background, i.e. spontaneous activity, and fixation conditions. The spectra vary little, except that the dip below 20 Hz becomes more prominent when spike rate increases. This can be explained by the potentially greater effect of a refractory period at higher spike rates. These spectra are normalized by spike rate.

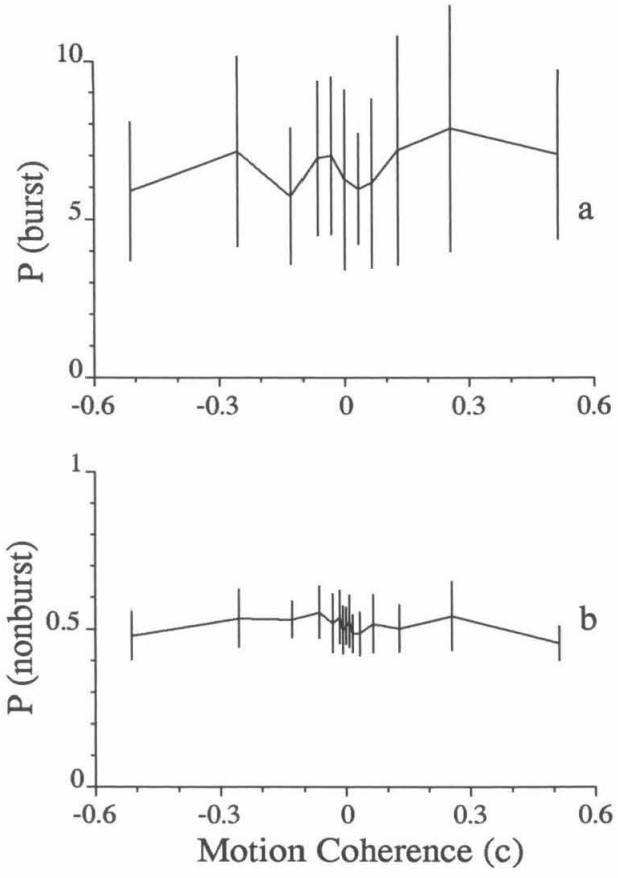

Figure 7.6: For most cells, P does not change substantially over the range from $c = 0$ to values well above the threshold level. The average value of P over all trials at each coherence level is plotted for a strongly bursting cell (top) and a strongly nonburst cell (bottom). Error bars show standard deviations.

this question, we computed the index of burstiness, P, for spike trains obtained under three different behavioral conditions. In the "choice" condition, the animals attended to the random dot stimuli with the intent of making a discrimination. The same visual stimuli were presented in the "fixation condition," but the animals were only rewarded for maintaining fixation on a visual target; no discrimination was required. Finally, the "background" refers to spontaneous neuronal activity that was acquired during the interval between trials.

The index of burstiness, P, did not vary between the "choice" and "fixation" conditions for the 82 cells for which "fixation" data was available (paired t-test, $p > 0.05$), suggesting that the monkey's intent to make a discrimination had no effect on the spectral peak. There was a significant change in P between the choice and background conditions (paired t-test, $p < 0.05$), but the effect was quite small: P decreased by an average of 3% for burst cells ($n = 122$) and increased by an average of 8% for nonburst cells ($n = 59$). We therefore infer that P is substantially independent of behavioral state, a conclusion that is supported by visual inspection of power spectra like those illustrated at the bottom of Fig. 7.5.

We next inquired whether the spectral peak developed or changed with time during the course of a two second trial. Analyzing the evolution of the shape of the power spectrum is difficult due to the small amount of data that most spike trains contain in a period as short as a few hundred milliseconds. The average spike rate over our entire database is 19 Hz (with a standard deviation of 18 Hz), so the exact placement of any one spike will have a large contribution toward the overall shape of the power spectrum for short windows. Because of this, we address a special case of this question which allows averaging over trials. This method is therefore limited to detecting changes which are locked to the stimulus onset.

We divided each trial (starting 336 msec after the onset of the stimulus to eliminate initial transients) into 6 equal time windows which overlapped by 1/3 of their width.

The average value of P was computed from individual spectra for all windows of similar time lags which fulfilled a minimum spike criterion of 5 spikes per window. Only 10% of burst cells and 21% of nonburst cells showed a significant correlation between P and time (Spearman rank-order correlation coefficient, $p < 0.05$). Of those cells, P increased by an average of 11% and 14% for burst and nonburst cells, respectively. Because P shows no correlation with time during the trial for 86% of cells and changes little for the other cells, we compute only one spectrum per trial in other analyses.

In a prior analysis of this data set, signal detection theory was used to compute a neuronal "threshold" that expressed the sensitivity of each neuron to motion signals in the display (Britten et al., 1992; see Methods). Threshold was defined to be the coherence value at which the neuron signaled the direction of motion with a criterion level of reliability. Thresholds varied widely among neurons in the data set, and we therefore tested for the hypothesis that burstiness as measured in the present analysis could be systematically related to the measure of sensitivity computed in the prior study. Fig. 7.7 shows cell threshold plotted against the index of burstiness, P, for all neurons that yielded a reliable estimate of P. The scatterplot contains no structure signifying a relationship between the two measures, an impression that is confirmed by calculation of a correlation coefficient ($r = 0.045$, p=0.61).

Psychophysical threshold also varied across these experiments since the testing conditions were changed to match the preferences of each cell (Britten et al., 1992). For some purposes, therefore, it is useful to express the sensitivity of each neuron relative to psychophysical sensitivity by calculating for each experiment the ratio of neuronal to psychophysical threshold. To determine whether P is related to cell sensitivity expressed in this manner, we calculated a correlation coefficient between the log of the "threshold ratio" and log P, but again we found no relationship ($r = 0.097$, p=0.28). Thus, the prominence of a peak in the 40 Hz region of the power

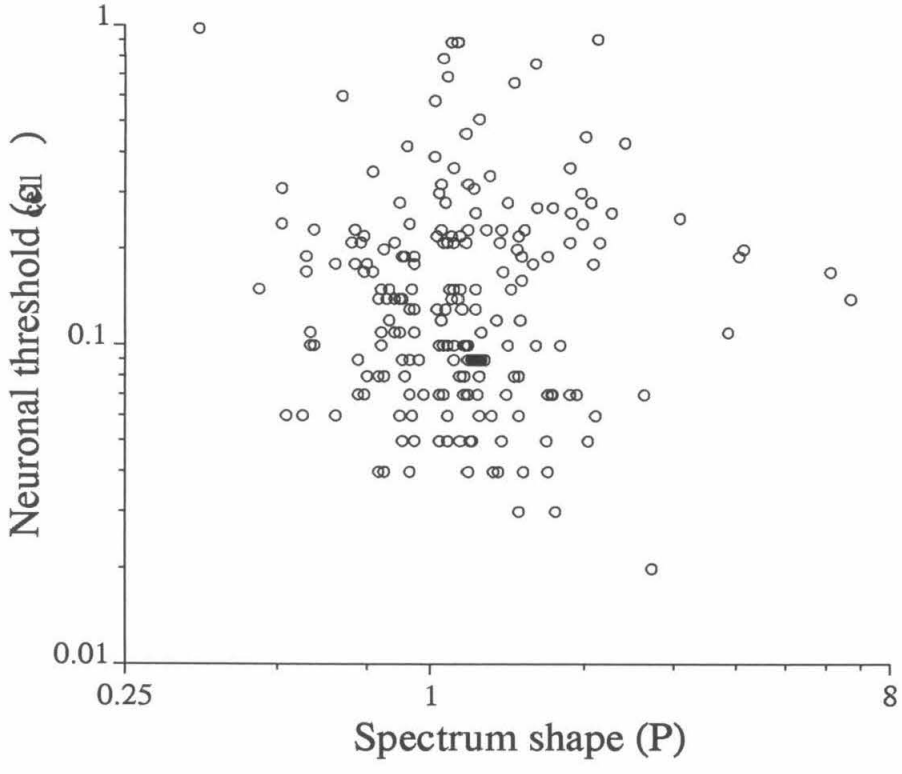

Figure 7.7: A comparison of neuronal threshold to the shape of the power spectrum. There is no significant correlation ($r = 0.045, p = 0.61$) between $\log(c_{cell})$ and $\log(P)$. Neuronal threshold, c_{cell}, is the coherence level which supports 82% correct decisions by an ideal observer counting total number of spikes. P is the ratio of the height of the peak (dip) in the 40 Hz band of $S'(f)$ to the baseline level for burst (nonburst) cells. We obtain similar results when correlating P against measures of the animal's psychophysical performance on motion discrimination (not shown).

spectrum does not correlate with prior measures of cell sensitivity.

Finally, we asked whether the prominence of the peak varied in a systematic way with the decision made by the monkey on successive presentations of a given motion condition. Previous analyses have demonstrated a trial-to-trial covariation between neuronal response and psychophysical decision when the response is considered to be the integrated spike count (Newsome et al., 1989b; Britten et al., 1995). We therefore wondered whether a measure of temporal structure in the spike train, P, would be similarly correlated with performance.

As in the prior analysis, we eliminated the influence of the visual stimulus itself on the monkey's decisions by carrying out the analysis only for $c = 0$ (completely random motion) and other small coherence levels for which neither the monkey nor the neuron discriminated the direction of motion at greater than chance levels (effectively random motion). On these trials, the monkey "guessed" the correct direction since the visual stimulus itself contained no effective information about the correct choice. For each cell we computed the difference between the average value of P for trials on which the monkey decided in favor of the neuron's preferred direction and the average value of P when the monkey decided in favor of the null direction. We found no statistically significant difference in P between these two conditions for either burst or nonburst cells (paired t-test, $p > 0.05$), and we therefore conclude that the prominence of the spectral peak in our data set and for our stimulus conditions is not related to the monkey's behavioral choice.

We found a similar result when testing whether P was correlated with correct versus incorrect decisions by the monkey at the coherence level closest to the monkey's psychophysical threshold, c_{system} (typically at $c = 0.128$). At this level, there are a significant number of incorrect response trials, and yet the monkey is not simply guessing. Again, we found no statistically significant difference (paired t-test, $p > 0.05$) for either burst or nonburst cells, and therefore conclude that the prominence

of the spectral peak is not related to correct and incorrect responses by the monkey.

Since it is a widely held belief that changes in temporal structure (such as an increase in burstiness) can result from cell damage caused by the electrode, we tested for a change in P from trial to trial over the course of the experiment. We found that 13% of burst cells showed a significant increase in P, 67% on average, during the experiment, while 12% showed a significant but small (only 8% on average) decrease in P. Among nonburst cells, 20% showed a significant increase, while 18% showed a significant decrease in P. The magnitudes of the increase and decrease among the nonburst cells were both 10% on average.

7.3.3 Treating Bursts as Events

We previously appealed to the neuronal threshold, c_{cell}, as a measure of an ideal observer's ability to decide the direction of motion of the stimulus based on the output of the neuron (Newsome et al., 1989a; Britten et al., 1992), assuming that the relevant neuronal output is the number of spikes fired during the stimulus period without considering whether those spikes occurred in bursts or as isolated action potentials. What happens if we quantify the neuron's output by the number of "events," where an event is either a burst or an isolated spike, and recompute an associated neuronal threshold, $\tilde{c}_{cell}$? One could well argue on biophysical grounds that a burst of spikes could be more powerful in evoking a postsynaptic response than the same number of isolated spikes.

Consistent with our definition of B, events are defined as the longest sequences of spikes with all interspike intervals less than or equal to 3 msec (values between 3 and 8 msec give very similar results). With this definition, single isolated spikes as well as bursts are counted as individual events. A nonburst cell will have nearly the same number of events as spikes, while a burst cell will have many fewer events

than spikes. Fig. 7.8a compares the tuning of a strongly bursting cell's response measured in spikes/sec (upper curve), events/sec (middle curve) and spikes/event (lower curve). Typical of our database, spikes/event is not tuned with c; therefore, the curve for events is similar to that for spikes, but scaled down by the average number of spikes per event, here 2.5. In Fig. 7.8b, the thick line shows the events/sec curve scaled up by 2.5 spikes/event so that it overlays the spikes/sec curve. From the relatively smaller standard deviations for normalized events/sec, it is clear that for this cell events/sec is a more useful neuronal signal for predicting the direction of coherent motion.

Because the neuronal code that carries motion information in cortex is not known, and since likewise we do not know whether neurons post-synaptic to the one recorded differentiate between bursts and isolated spikes, we tried various schemes for weighting the contribution of events to the output signal based on a function of the number of spikes per event. First, we weighted isolated spikes, i.e. single spike events, as 1 and bursts, events of multiple spikes, as α, with α varying between 0.5 and 8. We also used a different weighting scheme, where each event, irrespective of whether burst or isolated spike, is weighted according to its number of spikes raised to a power, β. Note that $\beta = 0$ corresponds to the first weighting scheme with $\alpha = 1$, and $\beta = 1$ corresponds to our original scheme which does not differentiate between bursts and isolated spikes. In addition, we consider $\beta = 1/2$ and $\beta = 2$.

To assess the advantage of these schemes, we recomputed neuronal thresholds based on the modified output signals for the 41 burst cells where the peak in the power spectrum was at least 50% above the baseline ($P \geq 1.5$). More weakly bursting cells are ignored because we expect no effect when isolated spikes greatly outnumber bursts. Fig. 7.9 shows the frequency histogram of $\tilde{c}_{cell}/c_{cell}$, where $\tilde{c}_{cell}$ is the neuronal threshold based on the modified signal. The shifts of the distributions are significant ($p < 0.05$) for all histograms shown except for $\alpha = 0.5$. Leftward shifts indicate

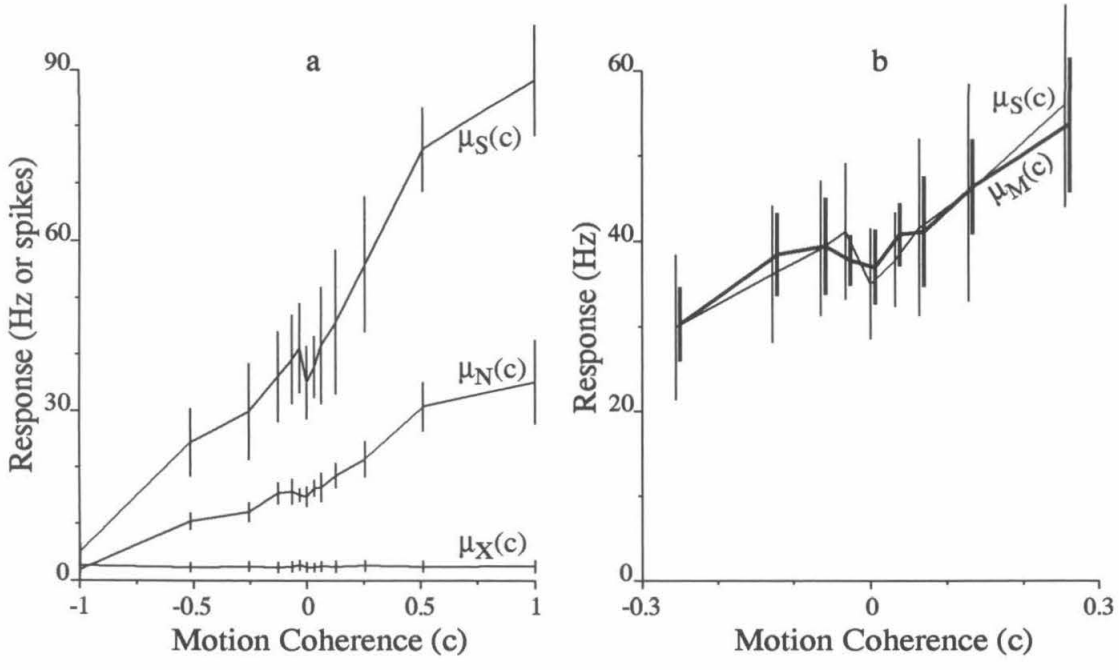

Figure 7.8: One example of how counting events rather than spikes can yield better direction discrimination. Events are defined to be either isolated spikes or bursts (groups of spikes with all intervals $\leq$ 3 msec). (a) The tuning of mean spike rate ($\mu_S(c)$), mean event rate ($\mu_N(c)$) and the number of spikes per event ($\mu_X(c)$) are plotted relative to motion coherence for the strongly bursting cell *j001*. $\mu_S(c)$ and $\mu_N(c)$, are tuned with c, but $\mu_X(c)$, is flat at about 2.5 and is treated as a constant, μ_X, in the text. (b) Focusing on the region around $c = 0$, $\mu_S(c)$ (thin line) is plotted against $\mu_M(c)$ (thick line, offset), computed by multiplying $\mu_N(c)$ from (a) by $\mu_X = 2.5$ spikes/event. It is apparent by the smaller standard deviations and similar slope that $\mu_M(c)$, and thus event rate, is a better basis for predicting c than is raw spike rate. Error bars show standard deviations.

that the thresholds improved (became lower) when the modified signal was used in place of spike count. The greatest improvement occurred for $\alpha = 1$ (i.e. $\beta = 0$) and corresponds to a 7.5% decrease in threshold. For three cells, thresholds were roughly cut in half. In other words, allowing an ideal observer to count bursts as single events enhances his ability to predict the direction of motion of the stimulus, on average.

Weighting bursts more ($\alpha > 1$) or less heavily ($\alpha = 0.5$) than single spikes did not improve thresholds. Squaring the number of spikes within the burst also lead to higher (worse) thresholds, while taking the square root yielded an improvement.

Based on these results, and on the relative variance-to-mean ratios for event count and spike count seen in Fig. 7.8, we believe that the improvement, particularly for $\alpha = 1$, is due to a reduction in relative variance which occurs by ignoring the number of spikes within events. This effect is easily demonstrated by a simple stochastic model. Consider the model that a bursty spike train is governed by two distributions, that of the number of events N and that of the number of spikes per event X. Assume that N is Poisson distributed with rate parameter $\mu_N(c)$, a function of stimulus coherence, and that X is distributed with mean μ_X and variance σ_X^2. Using basic results from the theory of branching processes (Feller, 1968), the mean of S, the number of spikes per trial, is then

$$\mu_S(c) = \mu_N(c)\mu_X \tag{7.5}$$

and the variance is (see Appendix for proof)

$$\sigma_S^2(c) = \mu_N(c)(\mu_X^2 + \sigma_X^2) \tag{7.6}$$

where we use the fact that $\mu_N(c)$ is both the mean and variance of the Poisson distribution for N. Rather than comparing the neuronal output signal S, based on spikes, directly to the event count N, we consider the random variable $M = \mu_X N$

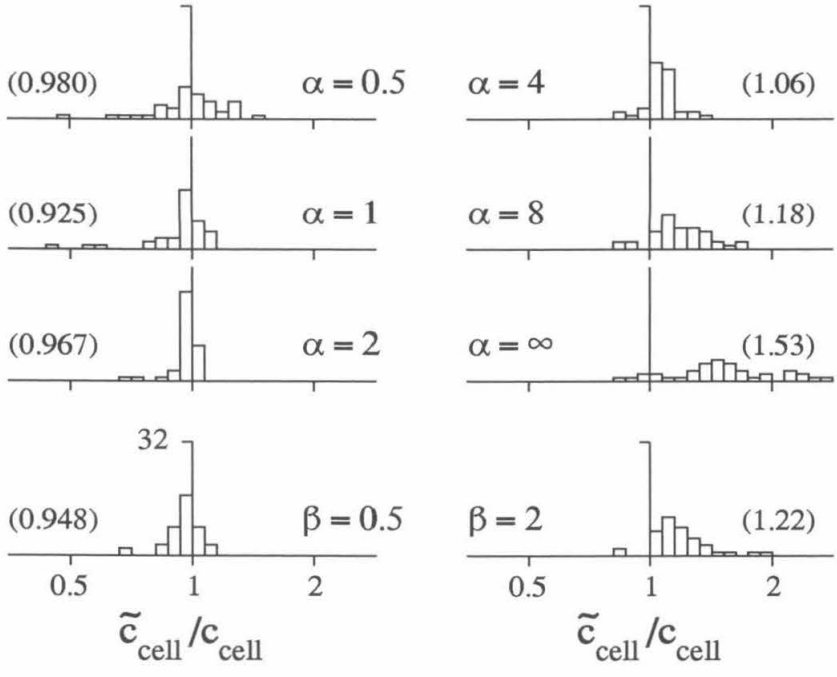

Figure 7.9: Comparing neuronal thresholds based on spike rate to those based on weighted event counts. For 41 burst cells ($P \geq 1.5$), frequency histograms show the threshold ratio $\tilde{c}_{cell}/c_{cell}$. In the upper six histograms, $\tilde{c}_{cell}$ is computed from ROC analysis based on the number of single spikes plus α times the number of bursts. The greatest leftward shift in the distribution (numbers in parentheses show means), representing the largest average improvement in neuronal threshold, is achieved for $\alpha = 1$, which corresponds to using $\mu_N(c)$, event count, as the neuronal signal. (The counts near 0.5 indicate cells for which the neuronal thresholds were roughly halved by this procedure.) For $2 \leq \alpha \leq 3$, this procedure is very similar to counting individual spikes, since bursts are composed of typically 2–3 spikes. Histograms for $\alpha = 0.75, 1.5, 3.0$ (not shown) have means $0.944, 0.940, 1.02$, respectively. As $\alpha \to \infty$, single spikes are ignored and only bursts are counted. The bottom histograms show results from two schemes which weight events based on the number of spikes/event raised to the power β. The square-root yields an improvement in threshold since it reduces the effect of variance in the number of spikes per event, while squaring emphasizes the variance, and worsens the thresholds. Overall, these plots indicate that an ideal observer with knowledge of the arrangement of spikes in bursts will be better able to predict the direction of motion, particularly at near-zero coherence levels, than an observer knowing only the total number of spikes.

which has the same mean as S, that is

$$\mu_M(c) = \mu_N(c)\mu_X. \tag{7.7}$$

This corresponds to multiplying the event/sec curve in Fig. 7.8a by 2.5 spikes/event so that it lies directly on top of the spikes/sec curve and allows direct comparison of the sizes of the standard deviations (i.e. Fig. 7.8b). The variance of the scaled-up event count M is given by

$$\sigma_M^2(c) = \mu_N(c)\mu_X^2, \tag{7.8}$$

which no longer has the term from eq. 7.6 involving the variance of the number of spikes within an event. Therefore, counting events rather than spikes should allow an observer to better predict the direction of motion of the stimulus under the conditions of this model. The critical condition here is that the number of spikes per event is independent of stimulus condition, and this appears to be the case, as demonstrated in Fig. 7.8, for most of the bursting cells, particularly at near-threshold coherence levels. This is further supported by our observation that P changes little with coherence level.

7.4 Stochastic Models

What stochastic models of neuronal firing can give rise to the observed power spectral densities and ISI distributions? To answer this question completely would require a detailed understanding of the biophysics of individual cortical cells as well as the dynamics and connectivity of the network in which the MT cells from which we recorded are embedded. Instead, we focus on the statistical properties of the discharge frequency of individual cells in a qualitative manner, bypassing the need for detailed single cell or network models. We believe that this method is justified by our results; we can model the power spectra and ISI distributions using very simple two- or four-

parameter distribution functions.

7.4.1 Poisson-Distributed Action Potentials

Our starting point is the fundamental result that the power spectrum of a random, Poisson train of action potentials is flat at all frequencies except for a delta peak at the origin. This agrees with our intuition that all spectral components should be equally represented in a completely random spike train. To be more precise, we will model a spike train as a shot-noise process $f(t)$, where the function $h(t)$ describes the shape of a single shot, here at first a single action potential, and $S_h(f)$ is the associated energy spectrum. A train of infinitely many action potentials is given by

$$g(t) = \sum_{i=-\infty}^{+\infty} h(t - t_i). \tag{7.9}$$

If the spikes occur at random, that is without any memory of the previous spike, but with an average firing rate λ, the power spectrum of such a random spike train is (Champeney, 1973)

$$S(f) = \lambda S_h(f) + 2\pi |\langle g(t) \rangle|^2 \delta(f), \tag{7.10}$$

where the average value of g is related to $h(t)$ by $\langle g(t) \rangle = \lambda \int h(t) dt$. In the familiar case where we assume that an individual action potential is adequately described by a delta impulse function, i.e. $h(t) = \delta(t)$, $S_h(f) = 1$, the above equation reduces to

$$S(f) = \lambda + 2\pi\lambda^2 \delta(f). \tag{7.11}$$

The interspike interval probability density function (ISI) for this case is given by

$$\mathrm{ISI}(\Delta t) = \lambda e^{-\lambda \Delta t} \qquad \Delta t > 0. \tag{7.12}$$

Thus, if we observe a Poisson spike train for a sufficiently long time, its interspike intervals should be distributed according to a single decaying exponential.

We will also make use of a more general expression for the power spectrum which includes cases where the occurrence of an action potential is dependent on the last time an action potential was initiated. We embody this dependency in the renewal density function $p(t)$. Assume that a spike was generated at time t_1. The probability for the next spike to occur between $t_1 + t$ and $t_1 + t + dt$ is given by $p(t)dt$ (for details, see Perkel et al., 1967). The power spectrum of shot-noise with this dependency among the "shots" is

$$S(f) = \lambda S_h(f) \left[1 + \int_{-\infty}^{\infty} e^{-i2\pi ft} p(t)\, dt \right], \tag{7.13}$$

for all values $f \neq 0$ (see Appendix L in Champeney, 1973). For a Poisson process, the probability of spiking per unit interval is always constant and is characterized by the mean rate; thus $p(t) = \lambda$.

Neurons, however, do not fire totally without memory, because for a variable time following the generation of an action potential the spiking threshold is elevated, making it more difficult to discharge the cell (absolute and relative refractory periods).

The effect of a refractory period can be modeled analytically with the help of the renewal density function. The shape of $p(t)$ can, indeed, be measured directly by computing the probability for the observed cell to fire an action potential in the short time interval $t_1 + t$ and $t_1 + t + dt$, assuming that it had fired at time t_1. For the binary data we have here (per sampling interval of $\Delta t = 1$ msec, either zero or one spike can occur), $p(t)$ is directly proportional to the auto-correlation function $R(t)$. For our non-bursting cells (e.g. cell **d**, Fig. 7.2), $R(t)$ (not shown) is well fitted by a constant minus a small Gaussian around the origin, indicating a reduced probability

of firing around $t = 0$. We therefore assume for the renewal density

$$p(t) = \lambda - \lambda e^{-\frac{t^2}{2\sigma^2}}. \tag{7.14}$$

Replacing $p(t)$ into eq. 7.13 yields the power spectrum of an infinite train of Poisson-distributed δ impulses with refractory period

$$S_{Poisson}(f) = \lambda(1 - \sqrt{2\pi}\lambda\sigma e^{-2(\pi f\sigma)^2}), \tag{7.15}$$

for $f \neq 0$. In order to ensure that $S_{Poisson}$ is always positive, the maximum firing rate must be limited: $\lambda \leq 1/(\sqrt{2\pi}\sigma)$. This spectrum, parameterized by two parameters, the mean rate λ and the width of the refractory period σ, is constant for large values of f but dips towards its minimum at $f = 0$. Fig. 7.10 shows $S_{Poisson}(f)$ for $\lambda = 40Hz$ and for $\sigma = 1, 2, 4, 8$ msec. A longer refractory period causes a deeper trough at low frequencies. Note that this result appears at odds with intuition, since a refractory period seems to demand a peak in the neighborhood of the inverse of the smallest interspike interval. However, this is only true if the firing rate λ is so high that the mean time between spikes approaches the refractory period. In that case, eq. 7.15 no longer describes the resulting spectrum. Clearly, in the case of a fixed absolute refractory period, the mean spike rate must be no greater than the inverse of the refractory period. The additional factor of $1/\sqrt{2\pi}$ arises because of the Gaussian depression in the renewal density used to model a stochastic refractory period.

We also numerically simulated this situation using Poisson generated shot-noise (with a mean spiking rate $\lambda = 58$ Hz) and a Gaussian distributed absolute refractory period (of 5 msec mean, 2 msec standard deviation and truncated below zero). In other words, each time a spike was generated, the Gaussian distribution specified the refractory period associated with that particular spike. After this refractory period,

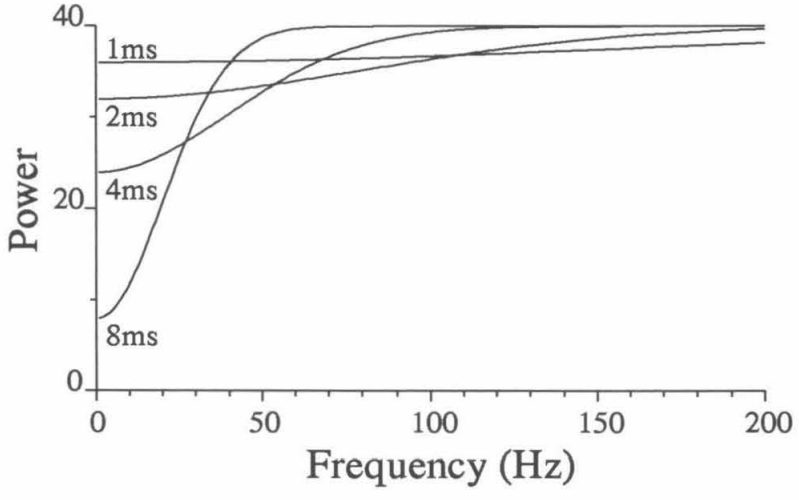

Figure 7.10: Derived power spectra $S_{Poisson}$ for a random (Poisson) spike train with refractory period. The power spectrum $S_{Poisson}$ for a Poisson process with a refractory period modeled as a Gaussian-shaped depression in the renewal density (see eq. 7.15) is plotted for $\sigma = 1, 2, 4, 8$ msec at $\lambda = 40$ Hz. As the length of the refractory period increases, the trough becomes deeper at lower frequencies. This model for $S(f)$ only holds if $\lambda \leq 1/(\sqrt{2\pi}\sigma)$. Outside of this range (i.e. if the inverse of the mean spike frequency is on the order of the refractory period) the spectrum develops a peak at that inverse of the mean refractory period, and the renewal density can no longer be modeled as a constant minus a Gaussian.

the probability for the next spike to occur is a constant $\lambda\Delta t$. We compare in Fig. 7.11 the spectrum and ISI distribution from an MT cell with a relatively flat spectrum and a dip at low frequency against this simple model. Both the synthetic and the experimentally recorded spike trains (see top portion of Fig. 7.11) are subject to the same analysis. The associated post-stimulus time histogram (PSTH) is flat for the computer-generated process since our model does not account for the transient component of the neuron's response. It is obvious that the ISI distribution and power spectrum for the synthetic process are very similar to those for the MT cell. In particular, both spectra show a dip at low frequencies.

We superimposed the analytical expression $S_{Poisson}(f)$ (with $\lambda = 58$ msec, $\sigma = 3.5$ msec) onto the power spectrum for the neuron (Fig. 7.11); it appears to provide an excellent fit to the computer-generated and the measured spectra. We performed this fitting procedure of $S'(f)$ against the two parameter function $S_{Poisson}(f)$ of eq. 7.15 for 61 nonburst MT cells. Six examples of the fits are shown in Fig. 7.12 for neurons with various firing rates, and the values of λ and σ for all nonburst cells are shown in Fig. 7.13. As expected, the refractory period shortens as firing frequency increases. Altogether, we find it remarkable that such a simple stochastic model of cell firing accounts for the shapes of the power spectra of many cells in our database.

7.4.2 Poisson-Distributed Bursts

Can we use a similar model to account for the 41 MT cells with a peak in their power spectrum whose amplitude is at least 50% higher than the baseline? To answer this question, we consider the interburst interval (IBI) distribution; if bursts occur at random but with a fixed absolute refractory period, their distribution should correspond to a shifted exponential, that is $\mathrm{IBI}(\Delta t) = \gamma e^{-\gamma(\Delta t - t_0)}$ for $t \geq t_0$ and 0 elsewhere, where t_0 is the duration of the absolute refractory period and γ is the mean rate for

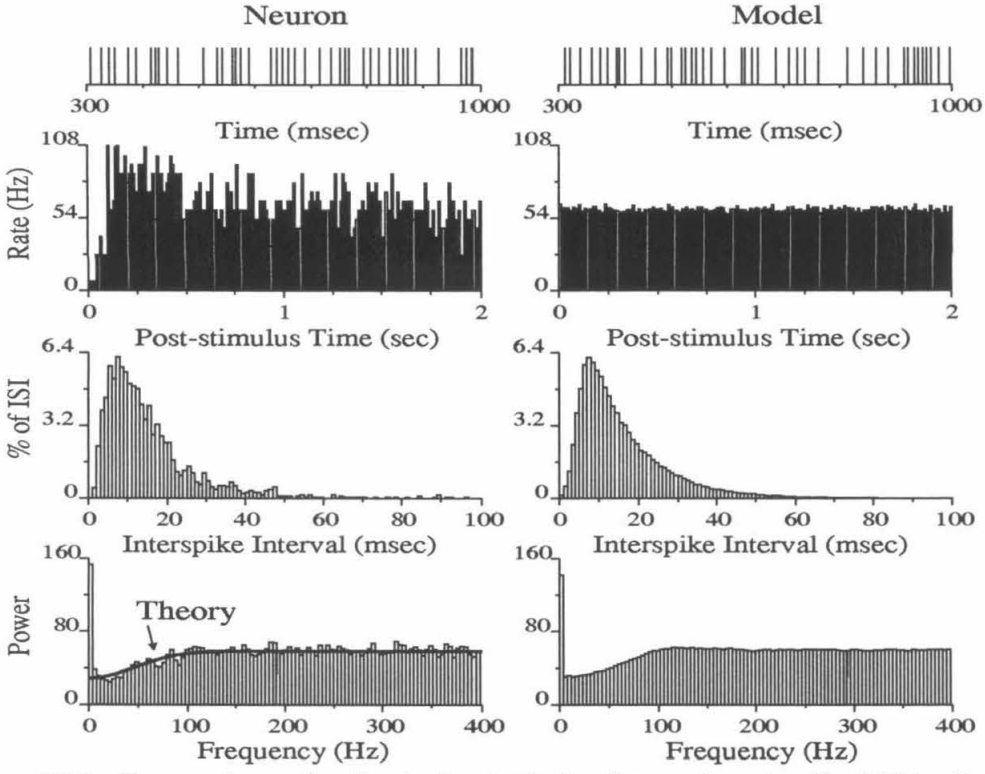

Figure 7.11: Comparison of spike train statistics for nonburst cell *e047* to those for a simple numerical model. The statistics for the neuron (left column) were computed by averaging over 15 trials at $c = 0.128$. The spikes (top trace) are less clustered than random, as demonstrated by the absence of short intervals in the ISI and the dip at low frequencies in $S'(f)$. The corresponding numerical model (right column) consists of a computer generated Poisson process (mean rate 86 Hz) superimposed with a Gaussian distributed refractory period (mean 5 msec, standard deviation 2.0 msec, truncated at 0 and renormalized). The model does not account for initial transients in the data and averages over the equivalent of 1000 two second long trials, so the PSTH is flat and all plots are less noisy for the model. The levels of PSTH and power spectrum (above 200 Hz) demonstrate that the resulting process has an overall mean rate matching that of the neuron. The absence of short intervals in the ISIs and the dips at low frequencies in the power spectra are in close agreement between the neuron and the model. This model is **not** intended to be a best fit for the data, but rather a demonstration that the location and size of the dip are qualitatively accounted for by a random process with a stochastic refractory period of appropriate duration. The solid curve superimposed on the neuron's spectrum (bottom left) corresponds to the analytical power spectrum for a Poisson process with a refractory period (eq. 7.15) with $\lambda = 58$ Hz and $\sigma = 3.5$ msec. Again, this qualitatively accounts for the dip.

Cell Name	λ (Hz)	σ (msec)
w014	80.3	2.0
j233	56.0	1.5
j200	46.8	1.0
e085	39.3	7.5
w213	30.6	6.5
j129	10.8	13.0

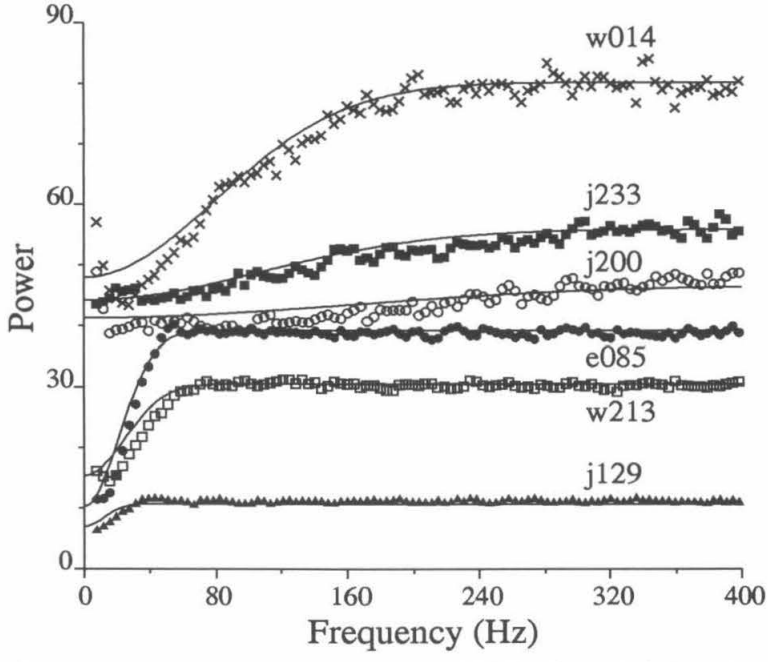

Figure 7.12: A comparison of power spectra, $S'(f)$, (points) for six nonbursting MT cells and their corresponding analytical curves, $S_{Poisson}(f)$, (solid lines) based on the expression for a process with randomly (Poisson) distributed spikes and a refractory period (eq. 7.15). These examples illustrate the ability of the simple analytical model to account qualitatively for the location and size of the dip in the power spectrum over a broad range of firing rates. The particular shape of the dip is often fit poorly since the form of the renewal density may not match the Gaussian shape imposed by analytical model (see eq. 7.14). The different levels of the various power function reflect the different mean spikes rates. The parameters of the analytical expressions are shown in the table.

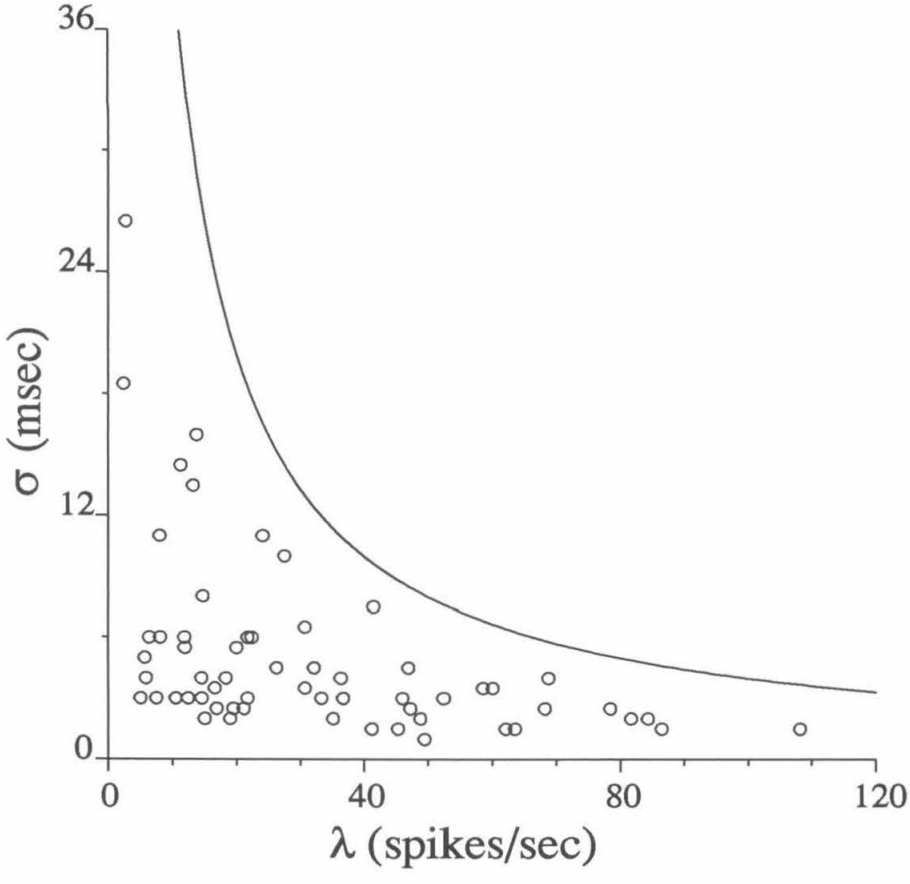

Figure 7.13: The distribution of parameters for fits of nonburst cells' spectra to the analytical model of the power spectrum of a Poisson process with a refractory period (see eq. 7.15). The refractory period parameter σ is plotted against the mean firing rate λ for 61 nonburst cells (points). The solid line shows the boundary outside of which the model no longer holds, i.e. for which $\lambda > 1/(\sqrt{2\pi}\sigma)$. To the upper right of this line, the firing rate becomes too high to support the corresponding refractory period under our model of the renewal density.

bursts. If, on the other hand, bursting cells are pacemakers, that is if they regularly fire in bursts at a fixed interval, the IBI should be sharply peaked around γ. Fig. 7.14 shows the average normalized IBI distribution for 37 cells (those with $P \geq 1.5$ and more than 200 interburst intervals). The logarithm of the distribution appears linear in the normalized (see Fig. 7.14 legend) time range of 40–160 msec and falls off at shorter intervals, consistent with a numerical model (thick curves) of Poisson distributed bursts with a burst-related refractory period which we develop in the next paragraph in terms of a single neuron.

To emulate the data for a single neuron shown in the left column of Fig. 7.15, we synthesize the following point process. Similar to the previous section, we generate "events" using a Poisson process (with $\lambda = 32$ Hz) with a Gaussian distributed refractory period (of mean 16 msec and 7 msec standard deviation; this distribution was truncated below zero and renormalized). Each event was then replaced with a burst of action potentials, i.e. δ-functions, where the length of the burst in milliseconds was approximately Gaussian distributed (mean 5.2, standard deviation 1.1 msec) and the spikes within the burst were chosen with approximately Gaussian spacing (mean 1.8, standard deviation 0.5 msec). The mean rate λ and the Gaussian refractory distribution were chosen to fit the measured IBI distribution. The parameters of the Gaussian distribution for the length of the burst and the density within the burst were also chosen to fit the neuronal data. If this model is simplified by assuming that the spikes within the bursts are generated by a Poisson process (similar to a model proposed by Smith and Smith, 1965), then the power spectrum would remain flat above 200 Hz, rather than gradually rising as seen at the bottom of Fig. 7.15.

The right side of Fig. 7.15 shows the resultant ISI and power spectrum, which are matched against similar functions for a bursting MT cell (left side of Fig. 7.15). What is surprising is that the synthetic data shows a peak in the power spectrum at about 31 Hz, **without** any underlying oscillations. How can this occur? A simple

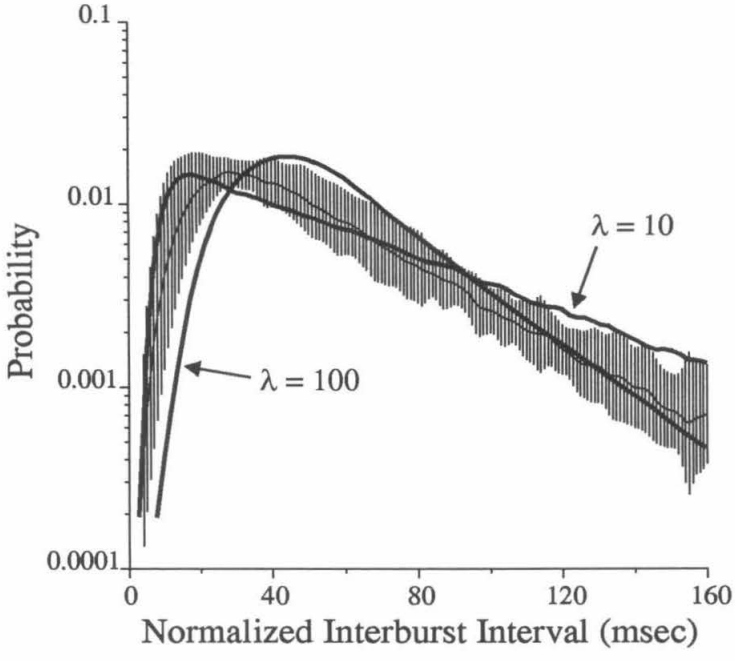

Figure 7.14: The average interburst interval (IBI) distribution for 37 burst cells ($P \geq 1.5$, 4 cells with too little data were discarded) and a numerical model. Individual IBIs were computed from all trials at $c = 0.512$ ($c = 0.256$ for 4 cells) using the 8 msec criterion (see text), normalized to have mean IBI equal to 60 msec (the population mean) by compressing or stretching the abscissa while preserving area, and smoothed with a Gaussian ($\sigma = 3.0$ msec) to eliminate isolated zero values. The average of logarithm of the IBI (thin line, error bars show standard deviation) is roughly linear from 40–160 msec. Since the mean IBI varied from 20–100 msec over the 37 cells, the linear range of 40–160 msec in the plot corresponds to ranges from 13–53 msec to 65–265 msec in actual time. Two thick curves show IBI distributions from a numerical model of a Poisson burst-generating process with Gaussian distributed burst-related refractory period ($\mu = 17$ msec, $\sigma = 8$ msec, as in Fig. 7.15, see text). Model curves are shown for $\lambda = 10$ and 100 bursts/sec. By varying λ, the model accounts for the slope and approximate shape of the normalized IBIs while holding constant the parameters of the stochastic refractory period. The variance of the data is smallest near the intersections of the model curves, consistent with the notion that varying the model parameters induces little change in the IBI in these regions. The model somewhat over-estimates the fraction of intervals in the linear range of 40–160 msec. Beyond 160 msec, values become undefined due to frequent zeros in the individual IBI histograms.

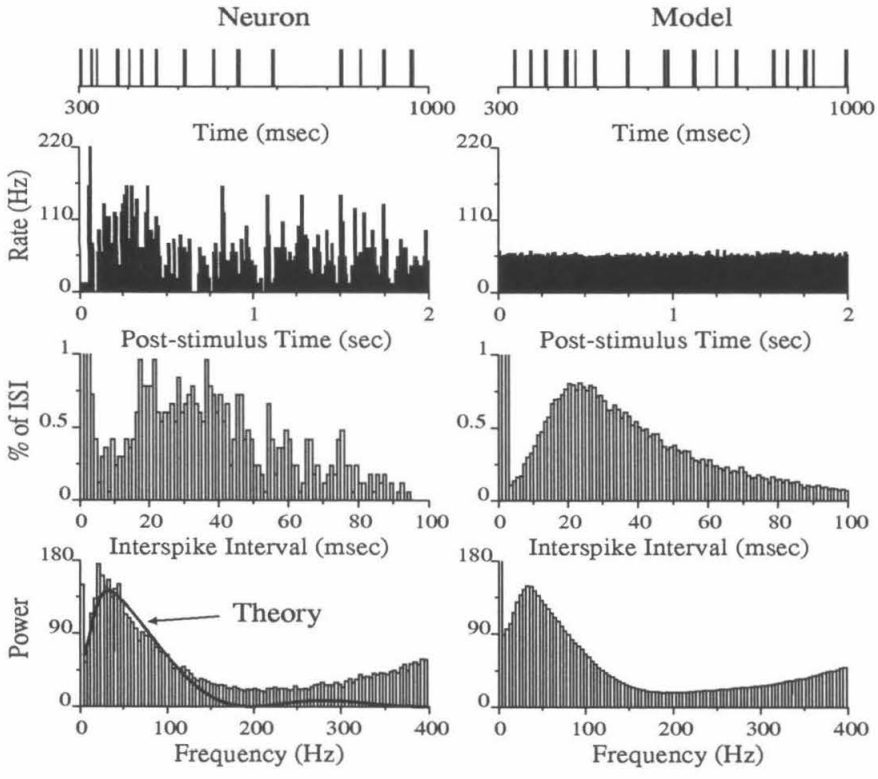

Figure 7.15: Comparison of spike train statistics for burst cell *j001* to those from a simple numerical model. The statistics for the neuron (left column) were computed by averaging over 15 trials at $c = 0.256$. The spikes (top trace) are more clustered than random, as demonstrated by the excess (62%, truncated on plot) of intervals in the 1, 2 and 3 msec bins of the ISI. The PSTH is particularly noisy because spikes occur in bursts. The corresponding numerical model (right column) is an extension of that used for nonburst cells (see Fig. 7.11) in which each spike generated from an underlying Poisson process with refractory period is now replaced by a burst of spikes where the burst length and the temporal structure of spikes within the burst are chosen to match the data (see text). Similar to the nonburst model, this is not intended to be a best fit to the data, but a demonstration that a process firing bursts randomly with a burst related refractory period can account qualitatively for the location, size, and shape of the peak near 33 Hz and the dips near the origin and at higher frequencies. The solid curve superimposed on the neuron's spectrum (bottom left) corresponds to the analytical power spectrum given by eq. 7.17 with $\lambda = 20.4$ events/sec, $\sigma = 12.5$ msec, $A = 550$ spikes/sec, and $L = 2.55$ msec. This cell is one of the 10% of cells whose peak in the 40 Hz range is twice as tall as the dip at higher frequencies.

analytical model proves to be insightful.

We again appeal to the power function of an infinite train of shot noise (where each individual shot is described by $h(t)$), with refractory period modeled with the renewal density $p(t)$ (eq. 7.13). While before we assumed that individual spikes can best be described using a $\delta(t)$ function, we now model a burst by a boxcar of amplitude A and half-width L centered around the origin. We set L to the half-width of the typical burst and A to normalize the area of the boxcar to account for the number of spikes within the typical burst. The energy spectrum associated with such an event is given by the square of a *sinc* function, i.e. by

$$S_h(\omega) = \frac{A^2}{\pi^2}\frac{\sin^2(2\pi Lf)}{f^2}. \tag{7.16}$$

The power spectrum of such Poisson events with a refractory period is

$$S_{burst}(f) = \lambda\frac{A^2}{\pi^2}\frac{\sin^2(2\pi Lf)}{f^2}(1 - \sqrt{2\pi}\lambda\sigma e^{-2(\pi f\sigma)^2}), \tag{7.17}$$

for $f \neq 0$.

We superimposed $S_{burst}(f)$ onto the neuron's spectrum in Fig. 7.15 (lower left) and find that both functions show a peak at the same frequency. The reason for the peak is the fact that S_{burst} is the product of $sinc(f)^2$, a decreasing function of f around the origin, and a monotonically increasing function, $1 - e^{-f^2}$. Fig. 7.16 shows the estimated power spectrum $S'(f)$ as well as the associated best-fit on the basis of eq. 7.17 for five burst cells. The analytical model does not account for variations in the burst width and occasional isolated spikes. Also, due to the use of the boxcar function to mimic bursts, we have no control over the fine structure of the spikes within the burst and therefore $S_{burst}(f)$ does not match well at high frequencies. What is important in this model is that the spectrum of this point process shows a peak, in

the absence of any underlying oscillator model. A similar result may be obtained in this case by using interburst intervals drawn from a broad Gaussian distribution with a mean value close to 25 msec. More neuronal data would be required to distinguish between the appropriateness of these models.

To emphasize the fact that the presence of bursts—in combination with a refractory period—can lead to a peak in the power spectrum, we used all 210 trials at different values of c for one particular cell, *j001*, and replaced every occurrence of a burst by a single spike, located at the center of the burst (Fig. 7.17). The associated power spectrum changes dramatically in character, from the usual peaked one to a flat spectrum with a dip at low frequency, compatible with the notion that once bursting has been accounted for, what remains are Poisson-distributed events modulated by the presence of a refractory period. If bursts would tend to occur every 25 msec or so, then this procedure should have led to a spectrum with a large peak around 40 Hz. For our data, bursts account satisfactorily for the peaks in the power spectrum. This is also witnessed by the rate of burst occurrence λ, shown in the table in Fig. 7.16, where λ is distributed between 10–20 Hz, below the range where the peaks appear in the corresponding spectra.

7.5 Discussion

The intent of the research reported here is to study some aspects of the time structure of spike trains recorded in cortex of the behaving monkey on the basis of the power spectral density, an order-dependent measure, and the interspike interval (ISI) histogram, an order-independent measure. Furthermore, we would like to assess whether knowing anything about the time-course of the cell's discharge can lead to more accurate predictions concerning the stimulus or the monkey's response than simply counting spikes. In this investigation, we have confined our analysis to data from

Cell Name	λ (Hz)	σ (msec)	A (Hz)	L (msec)
w150	13.0	22.4	558	2.3
w143	14.5	16.5	526	2.1
w190	17.7	17.5	421	2.0
w200	13.4	13.5	426	2.0
w180	15.4	11.8	374	1.8

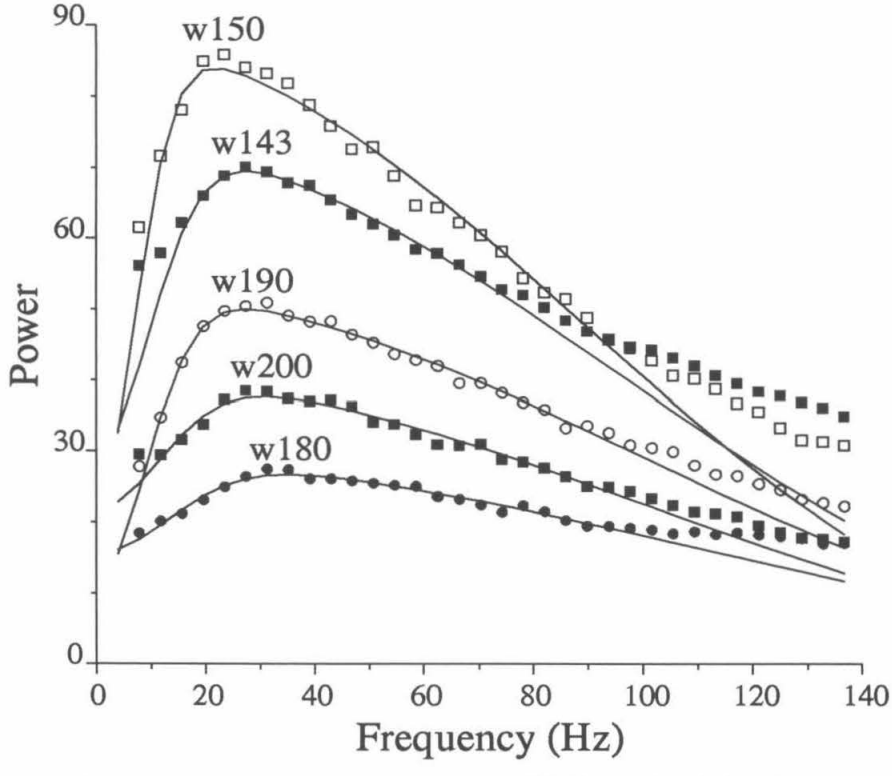

Figure 7.16: A comparison of power spectra, $S'(f)$, (points) for five strongly bursting MT cells and their corresponding analytical curves, $S_{burst}(f)$ (solid lines), based on the expression for a shot noise process with randomly (Poisson) distributed bursts (modeled as boxcar functions) with a burst-related refractory period (eq. 7.17). The parameters of the fits are shown in the table: λ is the mean rate of the shot noise, σ is the refractory period parameter, A is the height of the boxcar function (i.e. the spike rate within bursts), and L is the half-width of the boxcar function (i.e. the half-width of the burst). The analytical curves do not follow the spectrum at higher frequencies because they do not model individual actions potentials (see text).

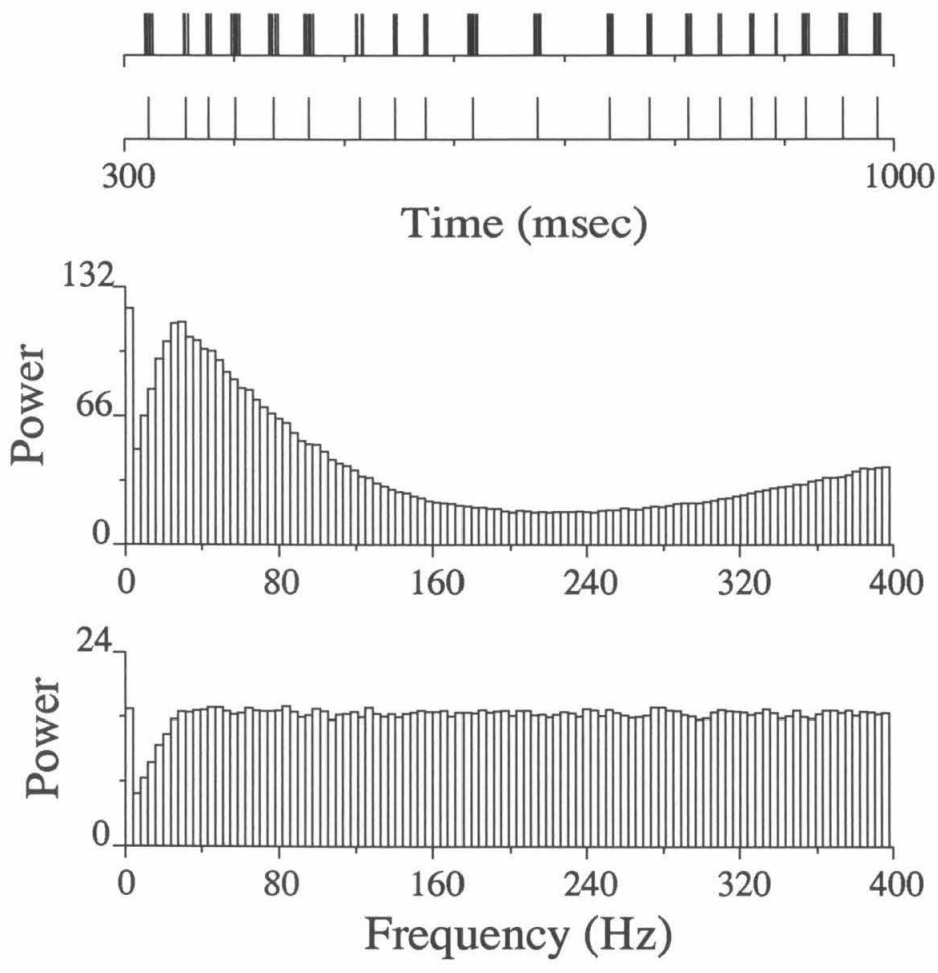

Figure 7.17: A demonstration that a peak in the power spectrum may be due to the presence of bursts, rather than regularity in their temporal alignment. When bursts from a neuron (upper impulse plot) are replaced by single spikes (lower impulse plot), the peak in the original power spectrum (upper spectrum) disappears (lower spectrum). A burst is defined as the longest subtrain of consecutive action potentials with no interspike interval greater than 8 msec (using 3 msec gives a very similar result). Each burst is considered to be an event at the mean occurrence time of all action potentials within that burst. This supports our notion that the burst themselves are randomly placed (with a burst-related refractory period) and are **not** locked to a regular oscillatory pattern.

well-isolated single units.

7.5.1 Random, Non-Bursting Cells

We found that about one-third (71) of our MT cells can be adequately described by a Poisson process of mean spiking rate λ with a refractory period (modeled here by a Gaussian distributed refractory period), in the sense that the experimentally determined power spectra $S'(f)$ and ISI distributions match the analytically (and numerically) determined ones (eq. 7.15 and Fig. 7.11). In particular, the power spectra are flat, with a dip towards low temporal frequencies. Spikes from these cells do not occur in bursts, that is they are almost always at least 4 msec (and usually much more) apart. We wish to point out that a Poisson process with refractory period is almost the simplest statistical description possible, with only two degrees of freedom. Yet, it appears to describe the measured discharge patterns relatively well. To our knowledge this is the first time that such a process is identified by its characteristic power spectrum.

In a related study (Softky and Koch, 1992, 1993), we computed the coefficient of variation C_V associated with the spike trains from the same data set of non-bursting MT cells (normalized for their non-stationary firing rates) as well as from V1 cells in the behaving monkey responding to bars and other textured stimuli (Knierim and Van Essen, 1992) and find values of $C_V \approx 1$, consistent with a Poisson process. We also analyzed the number of spikes occurring in different trials in response to the same stimulus and find that the variance in the number of spikes scales approximately as the average number of spikes to the 5/4 power (Softky and Koch, 1993). In a pure Poisson process, the variance should be equal to the mean, while for a fractal point process the variance can be larger than the mean (Teich, 1992; Usher et al., 1994). A number of studies have used this measure as indicative that the firing of cortical cells

in striate and extrastriate monkey cortex are consistent with a description of spiking as a Poisson process (Tolhurst et al., 1983; Parker and Hawken, 1985; Vogels et al., 1989; Zohary et al., 1990; Snowden et al., 1992). Thus, at least for long spike trains in the trained monkey, the associated ISI and the power spectrum are compatible with the notion that the underlying point process can be described by an almost memoryless Poisson process with refractory period.

7.5.2 Bursting Cells

More complex temporal dynamics are shown by the large fraction of cells (131 out of 212) that frequently discharge in bursts, i.e. 2–4 spikes within 2–8 msec or less (see the raster plots in the top row of Fig. 7.2). The fraction of the total number of spikes in a train which are less than 3.5 msec apart (our measure of "burstiness" B) ranges from an extreme value above 60% to 0. Unfortunately, we were not able to separate our 212 cells into two clearly segregated subpopulations using this or a related measure, since the distribution of cells varies continuously from strongly bursting to non-bursting (Fig. 7.3). Thus, any grouping of cells into "bursting" and "non-bursting" will be arbitrary to some extent. However, the amount of burstiness associated with individual cells remains relatively constant for all visual stimuli tested, and, in particular, is independent of motion coherence (Fig. 7.5, Fig. 7.6 and our definition of P). Furthermore, we found no systematic relationship between burstiness and the onset or the duration of the experiment for the majority of cells. Finally, it should be remembered that due to the perceptually demanding nature of the experiment, the monkey had to be highly alert throughout each trial. These observations argue against the possibility that the bursting is related to the onset of drowsiness or sleep.

Intracellular current injection into cells in rodent slices of sensory-motor cortex has revealed three distinct types of neurons (McCormick et al., 1985; Connors and

Gutnick, 1990; Agmon and Connors, 1992). The majority of these *in vitro* cells respond to the sustained current by a train of action potentials, which adapt within 50-100 msec to a more moderate discharge rate ("regular spiking" cells). A second class of neurons, only infrequently recorded from, is capable of high discharge rates with little or no adaptation ("fast spiking" cells). A third set of neurons respond to the depolarization by generating a short burst of 2-4 spikes, followed by a long hyperpolarization. This cycle of burst and hyperpolarization persists for as long as the current stimulus persists ("intrinsically bursting"). In slice tissue, the regular spiking cells correspond to pyramidal neurons, fast spiking cells to GABAergic non-spiny stellate cells and the intrinsically bursting cells to layer V pyramidal cells (Agmon and Connors, 1992). Because only little is known about the distribution of these cell types in the monkey, we can at present only speculate to what extent our "bursting" cells in MT correspond to these layer V intrinsically bursting cells characterized in slice preparations. However, the evidence presented above certainly suggests that the propensity of our MT cells to respond with bursts is not dependent on the nature of the visual stimulus, but rather appears to be an intrinsic property of certain cells.

What is the statistical distribution of bursts? We converted spike trains of some bursting neurons into "bursting trains" by the simple rule that a burst was defined as the longest subtrain with no interspike interval greater than 8 msec (using values as low as 3 msec here made only very little difference in the result). The resulting interburst interval distributions (IBI; see Fig. 7.14 for average) are not readily compatible with a neuronal process that generates bursts at any fixed temporal interval. In most cases, the IBI distributions can be fitted assuming a Poisson distribution of bursts combined with a Gaussian distributed burst refractory period (e.g. cell *j001*; see Figs. 7.14 and 7.15). In some cases, the IBI has a long tail, arguing against a simple exponential decay. We suspect that occasional excesses of long interburst intervals is partly the reason that the Poisson IBI model (thick curves, Fig. 7.14)

somewhat over-estimates the fraction of IBIs in the range of 40–160 msec. The relatively short duration of the trials considered here does not allow a conclusive study of intervals that fall beyond the 160 msec (normalized) value in Fig. 7.14. We never observed narrowly peaked IBIs, suggesting that individual MT cells are not acting as pacemakers. This is partly, however, a question of linguistic convention, since any cell having an IBI distribution with a single peak (such as that shown in Fig. 7.14), might *in principle* be considered to "oscillate" at the inverse of this peak. We do not, however, find this to be a very helpful definition.

7.5.3 Cells with a Peaked Spectrum

About two-thirds of all our MT cells (131 out of 212) have a single peak in their power spectra in the 40 Hz range (mean 41 Hz, standard deviation 9 Hz; top row in Fig. 7.2 and Fig. 7.15). Such a peak in the power spectrum is **not**, however, associated with ringing or oscillatory behavior in the Fourier transform of $S'(f)$, i.e. the autocorrelation function $R'(t)$. Only in a single cell (*j001*) during very high levels of motion coherence did we ever observe an oscillatory response in $R'(t)$. This appears to be quite different from the study of Kreiter and Singer (1992), who report that a large fraction of MT cells in the awake monkey show oscillations.

What simple statistical model of neuronal firing can give rise to a peak in the power spectrum? We show that the power spectrum of a Poisson process with a Gaussian distributed refractory period is monotonic increasing, leveling off towards a constant value at higher frequencies (eq. 7.15). The spectrum associated with a boxcar-like burst is $[sin(f)/f]^2$, a monotonic decreasing function around the origin. The power spectrum of a process which randomly fires bursts followed by a refractory period, is then given by the product of these two equations (eq. 7.17). Given that one function is increasing with f and the second one decreasing, the product of the two

will have a local maximum. In our case, if bursts are treated as boxcar functions that are about 4 msec wide, occur at a frequency of between 10–20 Hz, and are followed by a refractory period of between 10–25 msec (see table, Fig. 7.16), the peak in $S'(f)$ lies in the 20–50 Hz range (Fig. 7.15). These values were obtained from the distribution of the bursts themselves and can also be justified on biophysical grounds (Connors and Gutnick, 1990). We find it surprising to what extent simple analytical models can account for the observed interspike interval distributions and power spectra of bursting cells at frequencies less than 100 Hz. Our computer simulation of such a firing process, which differs from the analytical model by resolving the boxcar bursts into individual impulses and the renewal density into a stochastic refractory period, gives a better match of $S'(f)$ at higher frequencies.

If the occurrence of every burst in a spike train is replaced by a single action potential throughout the entire spike train (and isolated action potentials remain single spikes) the power spectrum $S'(f)$ totally changes its character (Fig. 7.17), from a spectrum with a peak to a flat spectrum with a dip at low frequencies, compatible with our notion that bursts themselves are distributed according to a Poisson distribution with a burst-related refractory period. If, for instance, the peak in the spectrum is due to periodically occurring bursts, our procedure should have revealed a spectrum with a clear peak, rather than the flat spectrum with a dip. We believe that our method of replacing bursts with "events" is a useful diagnostic tool for removing the confounding influence of bursts on the power spectrum, revealing the underlying dynamics.

Another way in which a neuronal "oscillator" model differs from our "random burst" model is in the distribution of interburst intervals; the former gives rise to an interburst interval distribution tightly clustered around the inverse of the oscillation period, while the latter model is associated with a decaying exponential IBI modified by a refractory period.

As discussed in the previous section, our data are consistent with the random burst model (Fig. 7.14); however, the two models are difficult to distinguish when the oscillator model becomes less regular. We can show using computer generated data what is expected in the case of the oscillator model. Fig. 7.18 demonstrates the appearance of the ISI and spectrum $S'(f)$ in the case of a neuron which fires isolated spikes (top) or bursts (bottom) with an approximately Gaussian ISI or IBI. In the case of isolated spikes, the power spectrum remains flat with a dip related to the apparent refractory period induced by the Gaussian ISI for distributions with a broad range of standard deviations ($\sigma \geq 12$ msec). Once the standard deviation becomes smaller (top Fig. 7.18, $\sigma = 7$ msec), a prominent peak arises in the spectrum (see arrow on figure) related to the inverse of the mean of the Gaussian ISI. The bottom portion of Fig. 7.18 shows a similar result, except that the isolated spikes are replaced by bursts (see the bimodal ISI distributions) and the peak due to small values of σ is superimposed on the peak due to bursts, explained earlier. The narrow, i.e. more localized, peak should be associated with oscillations since it arises due to the regularity of the ISI and **not** due to the interaction of bursting with the refractory period. Localized peaks in the power spectrum, although not observed in our data, are associated with ringing in the cross-correlogram which has been reported in data from other laboratories.

7.5.4 Treating Bursts as Signaling Events

Because of the possible special relevance of bursts to signaling in the brain (e.g. Koch and Crick, 1994), we evaluated to what extent bursts convey a different message from that conveyed by a collection of individual spikes. Following Cattaneo et al. (1981a,b) and Bonds (1992), we plotted the tuning curve as a function of motion coherence for three different measures of cell response (Fig. 8a): spikes per second, events per

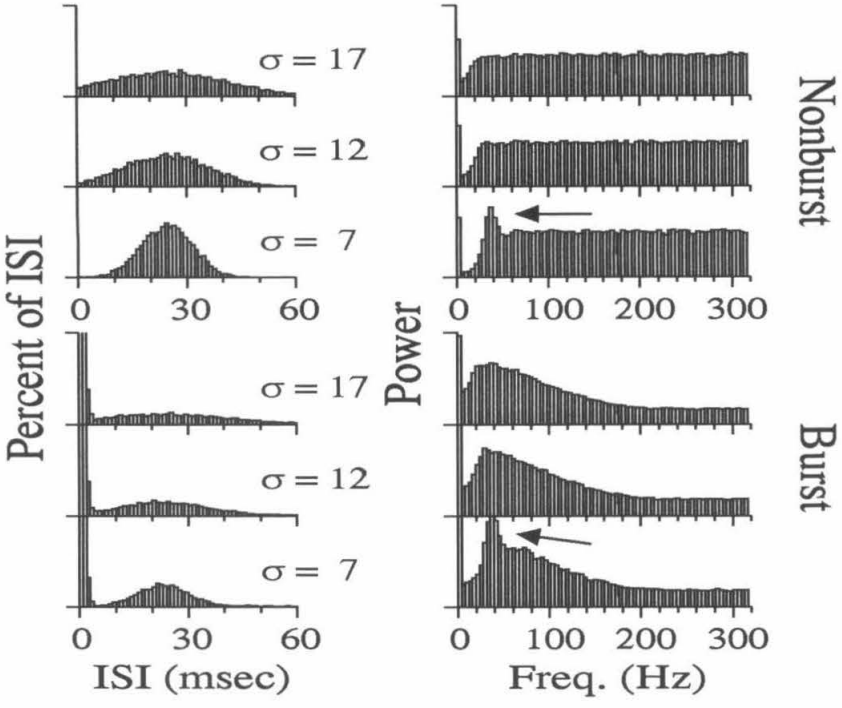

Figure 7.18: The difference between peaks in the power spectrum due to regular oscillation and peaks due to bursts. The top six plots show computer generated data for the Gaussian ISI model with mean 25 msec. The left side shows a series of ISI histograms in which the standard deviation, σ, of the Gaussian distribution is varied. For $\sigma > 12$ msec, there is no observable peak in the power spectra (right side), only a dip due to the apparent refractory effect imposed by the Gaussian. At $\sigma = 12$ msec, a peak is just beginning to form at 40 Hz (the inverse of the mean of the ISI), and as σ decreases further, the peak (shown by arrow for $\sigma = 7$ msec) rapidly increases in size. A similar progression occurs for the Gaussian IBI model (bottom set of six plots) in which spikes from the Gaussian ISI model are replaced by bursts. However, in the Gaussian IBI model, the spectrum has a broad bulge that sweeps up to a peak near 40 Hz due to the structure of the bursts (see Fig. 7.17), and this peak is present for all values of σ. A narrower peak appears above the broad peak due to bursts only when the variance of the Gaussian IBI becomes small. We therefore make a distinction between a peak in the power spectrum that is due to the local temporal structure of bursts and a peak that is due to a very narrow ISI or IBI distribution, which is a sign of regularity, i.e. oscillation, in the timing of spikes or bursts. It is common to see sharp peaks of the type pointed to by the arrows when neurons respond to artifacts in a stimulus, such as the 60 Hz refresh rate of a cathode-ray tube.

second—where an event is either a burst of spikes or an isolated spike—and spikes per event. Different from Bonds (1992), the average spikes per event (which includes individual spikes) does **not** vary with motion coherence c. Furthermore, if the event per second response is scaled up by the mean number of spikes per event, it closely follows the spike rate tuning curve (Fig. 7.8b). Thus, events or bursts *per se* have the same overall dependency on c as does the spike rate. However, due to elimination of the variability in the number of spikes per event when using event count, the variance of the scaled up curve, $\mu_M(c)$, is less than the variance associated with the mean spike rate, $\mu_S(c)$, in particular around low values of motion coherence.

Because of the reduced variability of this measure, we expect it to be a more reliable indicator of the direction of motion of the stimulus. This intuition is confirmed by our analysis. We repeated the original ROC analysis (based on signal detection theory) of Newsome et al. (1989a) but allowed more flexibility in defining the signal on which the ideal observer would operate. They assumed that an ideal observer (referred to as the "Stanford" observer) counts each action potential in the 2 sec long spike train from an MT cell and uses this number as the basis for his analysis. We added a twist to this by weighting bursts differently from single spikes. While each isolated action potential contributed 1 towards the final sum, all spikes in a burst make a total contribution of α. Setting $\alpha = 1$ implies that the entire burst only contributes as much as a single, isolated action potential towards the final decision.

The result, as shown in Fig. 7.9, is unambiguous. Setting $\alpha = 1$, corresponding to using the event rate as the neuronal signal, improves the neuronal threshold for most of the 41 strong burst cells we analyzed here. The mean improvement was 7.5%, and in 3 cells, the thresholds dropped by roughly a factor of 2. In other words, the "Caltech" ideal observer who distinguishes spikes from bursts can—for these cells—determine the correct direction of motion (using a statistical criterion) at a lower level of coherence c than the "Stanford" observer. $\alpha = 0.5$ and 2 gave smaller

improvements (as did $\alpha = 0.75$ and 1.5; not shown). As α is made larger ($\alpha \to \infty$), bursts are more and more emphasized at the expense of single isolated spikes and the thresholds increase by 53%. The threshold also increases when bursts are weighted by the square of the number of spikes in each burst ($\beta = 2$). Weighting events by the square-root of the number of spikes per event, ($\beta = 0.5$) on the other hand, decreases thresholds (Fig. 7.9), since it decreases the variability contributed to the final signal.

From the point of view of our fictitious pair of ideal observers, the Caltech observer does better than his Stanford counterpart by replacing each occurrence of a burst of spikes by a single spike. In this sense it can be said that a crude measure of the temporal organization of spike trains does better in terms of signal detection theory than a simple spike count. This is not to say, however, that more sophisticated measures of temporal organization, possibly taking account of the simultaneous activity of many neurons, cannot do better yet (e.g. Richmond and Optican, 1992; Aertsen et al., 1989; Singer, 1994).

We do not know at this point the code that neurons postsynaptic to MT cells use to decide whether the stimuli move in one or the other direction. The fact that the neuronal threshold of many cells is frequently lower than the psychophysical threshold of the entire animal (Newsome et al., 1989a) requires an explanation as to why the animal does not do better than it does (invoking correlated activity among cells and population coding; Britten et al., 1992). Our results point to an additional explanation: if bursts are substantially more efficient in elevating postsynaptic firing rate than isolated spikes, corresponding to $\alpha \gg 1$, thresholds would increase and averaging over many cells would be required in order to mimic the psychophysical thresholds. It is important at some point that this question be resolved experimentally, possibly using a combination of *in vivo* slice techniques with behavioral studies.

7.5.5 Functional Considerations

What is the function of bursts? Why should cortex have two types of long-range projection cells, one signaling isolated spikes and the other responding frequently with bursts of spikes? It has been argued (Koch and Crick, 1994) that bursting neurons are much more efficient at accumulating calcium in their axonal terminals than cells that fire isolated spikes (that is, four spikes within a 10 msec interval cause a much larger increase in intracellular calcium at the end of the last spike than four spikes within a 40 msec interval). Because intracellular calcium accumulation in the presynaptic terminal is thought to be mainly responsible for various forms of short-term potentiation (in particular facilitation and augmentation; Magleby, 1987), it may well be that the primary function of bursting neurons is to induce this non-Hebbian (that is, non-associative) type of synaptic plasticity at its postsynaptic targets outside of the cortical system. In essence, the burst of spikes could turn on short-term memory which would then decay over several seconds (see also Crick, 1984). One might then expect there to be a relationship between bursting and short-term learning.

It is important to know whether our "bursting" cells correspond to the "intrinsically bursting" cells identified by intracellular current injection. The latter cells appear to be confined (at least in rat and guinea pig slice) to layer V (Agmon and Connors, 1992). In rat area 17, these cells have been shown to project outside cortex, in this case to the ipsilateral superior colliculus, while the remaining pyramidal cells in layer V project to the contralateral striate cortex (Kasper et al., 1991). Recent *in vivo* recordings in awake cat motor cortex has revealed that cells at or below a depth of 800 μm (corresponding to layer 5) show strong bursting activity (Baranyi et al., 1993). It is not known to what extent such cells exist or are localized to particular layers in primate cortex.

7.5.6 Correlation to Behavior

Finally, we return to a question that provided primary motivation for this study. Is the animal's perception of the stimulus, as evidenced by performance, influenced by temporal structure in the spike trains, particularly with respect to the peak in the power spectrum near 40 Hz? We correlated the presence and strength of the peak in the spectrum to both the stimulus and the behavior of the monkey. As witnessed by Figs. 7.5 and 7.6, we found no significant correlation between the fraction of dots moving in one or the other direction and P, the measure of the peak in the power spectrum, for most cells. We repeated this measurement using the integral of power in the 40 Hz band of $S'(f)$ with similar results. Furthermore, if the monkey is not forced to respond to the stimulus or even in the absence of the motion stimulus, the basic propensity of a cell to show this peak remains. This is related, of course, to our earlier result that bursting in these cells does not depend on stimulus conditions.

We find a similar lack of correlation between P and the various measures of behavior used in a previous comparison of neuronal responses and psychophysical performance (Newsome et al., 1989a; Britten et al., 1992). For instance, we tested for significant changes in the distribution of P when the monkey made correct vs. incorrect decisions at near-threshold coherence levels and when the monkey made preferred vs. null guesses for zero coherence motion, but we found no significant (paired t-test, $p > 0.05$) correlations.

Fig. 7.7 shows another attempt at studying the relationship between the peak in the spectrum and the behavior of the monkey. As is evident, no correlation exists between P and the level of the neuronal threshold, c_{cell}, that is the fraction of dots moving in the cells preferred direction at which the cell can "decide" the correct direction of motion (using an ROC criterion; Newsome et al., 1989a; Britten et al., 1992). A similar lack of correlation exists between P and the ratio of single cell

thresholds to the threshold of the animal (c_{cell}/c_{system}) and between P and the decision related probability of each neuron (not shown). Thus, for our stimulus conditions, the presence or strength of a peak in the power spectrum of well-isolated units does not tell us anything about the behavior of these animals.

When analyzing our data set, it should be kept in mind that the three monkeys from which the MT cells were recorded were extensively trained using operant conditioning techniques. In order to perform correctly the motion discrimination task at threshold levels, up to 6 months of training were required (Britten et al., 1992). We analyzed in a preliminary manner data from MT cells from one naive monkey who was only trained to fixate (Zohary and Bair, personal communication), and we found no significant difference in the distribution of burst and nonburst cells or in the shape of the associated power spectra. However, it is possible that more subtle differences in the temporal fine-structure exists between cells in the naive and in the trained animal.

Although we cannot say how well our results will generalize to other cortical areas or other behavioral tasks, we believe that the approach taken here—correlating temporal structure in spike trains to the simultaneous behavior of an awake animal—is a necessary step in establishing the role played by temporal firing patterns in the animal's perception of visual stimulation.

7.6 Appendix

Here we assume that a cell fires N events during the course of a fixed duration stimulus and that the i^{th} event is composed of X_i spikes, where N and X_i are independent random variables, $1 < i < N$. All of the X_i are mutually independent and drawn from a common distribution, and we use X to refer to a variable drawn from this distribution with no reference to a particular event. The total number of spikes fired

during the trial is

$$S_N = X_1 + X_2 + \cdots + X_N \tag{7.18}$$

and we will prove for the random variable S_N that

$$\begin{aligned} \mathrm{E}(S_N) &= \mathrm{E}(N)\mathrm{E}(X) && (7.19) \\ \mathrm{VAR}(S_N) &= \mathrm{E}(N)\mathrm{VAR}(X) + \mathrm{VAR}(N)\mathrm{E}^2(X). && (7.20) \end{aligned}$$

We use E for expectation, VAR for variance and P for probability.

The probability that k spikes are fired during a trial is

$$\mathrm{P}\{S_N = k\} = \sum_{n=0}^{\infty} \mathrm{P}\{S_N = k | N = n\}\mathrm{P}\{N = n\}. \tag{7.21}$$

Using this and the definition of expectation,

$$\begin{aligned} \mathrm{E}(S_N) &= \sum_{k=0}^{\infty} k\mathrm{P}\{S_N = k\} && (7.22) \\ &= \sum_{k=0}^{\infty} k \sum_{n=0}^{\infty} \mathrm{P}\{S_N = k | N = n\}\mathrm{P}\{N = n\} && (7.23) \\ &= \sum_{n=0}^{\infty} \mathrm{P}\{N = n\} \sum_{k=0}^{\infty} k\mathrm{P}\{S_N = k | N = n\} && (7.24) \\ &= \sum_{n=0}^{\infty} \mathrm{P}\{N = n\}\mathrm{E}(S_N | N = n). && (7.25) \end{aligned}$$

Evaluating the conditional expectation in the previous line, we get

$$\mathrm{E}(S_N | N = n) = \mathrm{E}\left(\sum_{i=1}^{n} X_i\right) = n\mathrm{E}(X), \tag{7.26}$$

which, when substituted into eq. 7.25, completes the proof of eq. 7.19

$$\mathrm{E}(S_N) = \mathrm{E}(X) \sum_{n=0}^{\infty} n\mathrm{P}\{N = n\} = \mathrm{E}(X)\mathrm{E}(Y). \tag{7.27}$$

To compute the variance of the number of spikes, we first compute $\mathrm{E}(S_N^2)$. Reasoning as in eqs. 7.22–7.25 above, but now for S_N^2 rather than S_N, we get

$$\mathrm{E}(S_N^2) = \sum_{n=0}^{\infty} \mathrm{P}\{N = n\}\mathrm{E}(S_N^2 | N = n). \tag{7.28}$$

Evaluating the conditional expectation from the previous line, we get

$$\begin{aligned} \mathrm{E}(S_N^2 | N = n) = \mathrm{E}\left(\sum_{i=1}^{n} X_i\right)^2 &= \mathrm{E}\left(\sum_{i=1}^{n} X_i^2 + \sum_{i=1}^{n}\sum_{j\neq i} X_i X_j\right) \\ &= n\mathrm{E}(X^2) + (n^2 - n)\mathrm{E}^2(X), \end{aligned}$$

which on substitution into eq. 7.28 yields

$$\begin{aligned} \mathrm{E}(S_N^2) &= \mathrm{E}(X^2) \sum_{n=0}^{\infty} n\mathrm{P}\{N = n\} + \mathrm{E}^2(X) \sum_{n=0}^{\infty} n^2\mathrm{P}\{N = n\} - \mathrm{E}^2(X) \sum_{n=0}^{\infty} n\mathrm{P}\{N = n\} \\ &= \mathrm{E}(X^2)\mathrm{E}(N) + \mathrm{E}^2(X)\mathrm{E}(N^2) - \mathrm{E}^2(X)\mathrm{E}(N) \\ &= \mathrm{E}(N)\mathrm{VAR}(X) + \mathrm{E}^2(X)\mathrm{E}(N^2). \end{aligned}$$

Using the expressions for $\mathrm{E}(S_N)$ and $\mathrm{E}(S_N^2)$, the variance of the number of spikes is

$$\begin{aligned} \mathrm{VAR}(S_N) &= \mathrm{E}(S_N^2) - \mathrm{E}^2(S_N) \\ &= \mathrm{E}(N)\mathrm{VAR}(X) + \mathrm{E}^2(X)\mathrm{E}(N^2) - \mathrm{E}^2(X)\mathrm{E}^2(N) \\ &= \mathrm{E}(N)\mathrm{VAR}(X) + \mathrm{E}^2(X)\mathrm{VAR}(N), \end{aligned}$$

which completes the verification of eq. 7.20.

In eq. 7.6, E(N) and VAR(N) are replaced by the mean (which is equal to the variance for a Poisson distribution) number of events, $\mu_N(c)$, and the mean and variance of the number of spikes per event are μ_X and $\sigma^2(X)$.

Chapter 8

Effect of the Refractory Period on Power Spectrum

This chapter summarizes an analytical treatment of the effect of the refractory period on the power spectrum that begins directly from the inter-spike interval distribution, unlike the analytical results in Chapter 7 which assumed that the renewal density was known. The work here is based on derivations by Prof. Joel N. Franklin in the Applied Mathematics Department at Caltech. The full mathematical derivations are given in Franklin and Bair (1995).

The interspike intervals in steady-state neuron firing are assumed to be independent, identically distributed random variables. In the simplest model discussed, each interval is assumed to be the sum of a random neuron refractory period and a statistically independent interval due to a stationary external process, whose statistics are assumed known. The power spectral density (hence the autocorrelation) of the composite neuron-firing renewal process is derived from the known spectrum of the external process and from the unknown spectrum of the neuron-refraction process. Two models are demonstrated that may produce peaks in the power spectrum near

40 Hz. The results are applied to the MT database.

8.1 Introduction

There has been recent interest in the use of the Fourier power spectrum for analyzing temporal structure in trains of action potentials recorded from neurons (Ghose and Freeman, 1992; Bair et al., 1994). This is largely due to the reports of stimulus-induced oscillation near 40 Hz in neural activity recorded in the visual cortex of the anesthetized cat (Eckhorn et al., 1988; Gray and Singer, 1989; Gray et al., 1989; Ghose and Freeman, 1992) and the awake monkey (Kreiter and Singer, 1992).

In Chapter 7 we found that most of the spike trains are well modeled by a Poisson shot-noise process modulated by a refractory period where the shots are either δ-functions or boxcar functions, representing individual action potentials or bursts (temporal clusters) of action potentials, respectively. In both cases, the refractory period is implemented in the model using a Gaussian depression in the renewal density function (see Perkel et al., 1967) for the shot-noise process, and the results of Champeney (Champeney, 1973) are used to compute the resulting power spectrum.

In the present work, instead of relying on the renewal density to model the refractory period, we begin with a function that is explicitly designed to model the absolute and relative refractory period of the neuron. This density function describes the neuron's stochastic dead time following an action potential and is the first of two contributions to the period between successive action potentials. The remaining contribution depends on the underlying process describing the input to the cell. We allow this input process to be any stationary renewal process, not just the Poisson process used in Chapter 7. Using the refraction function and the probability density function for the intervals of the underlying process, we can derive the power spectrum.

The results allow quantitative comparisons of neuronal spike trains to a variety

of stochastic models. In particular, they enable us to determine how much of the regularity in successive interspike intervals in a spike train can be explained by a renewal process with a refractory period or to what extent it is necessary to assume some underlying oscillatory mechanism to explain regularity. In addition, because our model distinguishes between the contribution to the interspike interval from refraction and the contribution from input to the neuron, we believe that it provides a framework for separately characterizing a neuron's refractory period and providing a stochastic description of its inputs.

We apply the mathematical results to neuronal spike trains recorded in a previous study (Britten et al., 1992) to demonstrate four points. (1) A fixed stochastic description of a refractory period is adequate to describe the output of some neurons even when the input stimulation and the output firing rate vary dramatically. (2) The assumption of a renewal process for the generation of interspike intervals is consistent with data. (3) Spike patterns such as bursts are accounted for by the theory and result in predictable peaks in the power spectra of spike trains that are not due to regular timing. (4) If spike rate (or burst rate) becomes fast relative to the refractory period duration, peaks may develop in the power spectrum near the inverse of the peak in the interval density, and these peaks are the result of regular timing.

Many others have studied the effect of refraction on neuronal spike train statistics. Teich et al. (1978) considered the model of a Poisson process modified by a variable dead time which is very similar to that used in Chapter 7 and which is encompassed in the theory developed here. Teich and Diament (1980) modeled the relative refractory period using a gradual recovery function. Related mathematical models concerning the power spectra of impulse processes appear in Lukes (1961), Beutler and Leneman (1968), and Heiden (1969).

8.2 Mathematical Model

We outline the model and summarize the mathematical results. This model will be applied to the MT data in Section 8.3. For complete mathematical derivations, see Franklin and Bair (1995).

We will analyze the steady-state firing of a single neuron. We shall suppose that the output potential is a renewal process

$$x(t) = \sum_{k=-\infty}^{+\infty} h(t - t_k), \tag{8.1}$$

where the t_k are the successive firing instants and where $h(t - t_k)$ is the output of the single firing at the instant t_k. The function $h(t)$ is given. The ensemble $\{t_k\}$ is random. The successive positive lags $\tau_k = t_k - t_{k-1}$ are assumed to be statistically independent with the same p.d.f. (probability-density function) $f(\tau)$ for $\tau > 0$. Although the successive lags τ_k are random, their p.d.f. $f(\tau)$ is assumed to be known.

We suppose that $f(\tau)$ is continuous for $\tau \geq 0$ except for isolated simple jump discontinuities, and we will assume that $f(\tau)$ has finite moments

$$\int_0^\infty \tau^k f(\tau)\, dt \tag{8.2}$$

for $k = 0, 1, 2$, and preferably also for $k = 3$ and 4. We assume that the wave form $h(t)$ has a finite Fourier transform

$$H(\omega) = \int_{-\infty}^{\infty} h(t) e^{-i\omega t}\, dt, \qquad (-\infty < \omega < \infty). \tag{8.3}$$

We do allow $h(t)$ to be the Dirac δ-function. We expect, but do not require, that $h(t) \equiv 0$ for $t < 0$.

To study neuron refraction, we shall hypothesize a simple model for the successive

lags $\tau_k = t_k - t_{k-1}$. We shall suppose that each lag τ_k is the sum of two parts

$$\tau_k = \tau_{k_1} + \tau_{k_2} \tag{8.4}$$

where τ_{k_1} arises from refraction, and where τ_{k_2} arises from all other causes, for example from an assumed random input potential.

We shall suppose that the parts τ_{k_1} and τ_{k_2} are statistically independent, with probability densities $f_1(\tau)$ and $f_2(\tau)$ that are independent of k. Thus, the sum τ has the p.d.f.

$$f(\tau) = f_1(\tau) * f_2(\tau) = \int_0^\tau f_1(\tau - \lambda) f_2(\lambda)\, d\lambda. \tag{8.5}$$

The ensemble of random variables $\{\tau_{k_j}\}$ is supposed to be statistically independent for $-\infty < k < \infty$ and $j = 1, 2$.

We assume that the random variables τ_{k_1} and τ_{k_2} have finite means μ_1, μ_2 and finite variances σ_1^2 and σ_2^2. Then the total lags $\tau_k = \tau_{k_1} + \tau_{k_2}$ have the common mean $\mu = \mu_1 + \mu_2$ and the common variance $\sigma^2 = \sigma_1^2 + \sigma_2^2$.

We will show how the mean μ and the variance σ^2 appear in the power spectral density of the steady-state neuron output potential $x(t)$.

The power spectral density $S_x(\omega)$ is defined as the Fourier transform of the autocorrelation $R_x(\tau)$:

$$S_x(\omega) = \int_{-\infty}^{\infty} R_x(\tau) e^{-i\omega\tau}\, d\tau, \tag{8.6}$$

where the autocorrelation $R_x(\tau)$ is defined as the expected value

$$R_x(\tau) = \mathrm{E}\left[x(t)x(t-\tau)\right], \qquad (-\infty < \tau < \infty). \tag{8.7}$$

The autocorrelation is independent of t, by the assumption that the ensemble of firing instants $\{t_k\}$ is a stationary random process.

It can be shown (Franklin and Bair, 1995) that $x(t)$ has the power spectral density

$$S_x(\omega) = |H(\omega)|^2 \nu \left\{ 2\pi\nu\delta(\omega) + 1 + 2\mathrm{Re}\, \frac{\hat{f}(i\omega)}{1 - \hat{f}(i\omega)} \right\} \tag{8.8}$$

where ν is the average firing frequency $\nu = \mu^{-1}$, and where $\hat{f}(s)$ is the Laplace transform of the lag p.d.f:

$$\hat{f}(s) = \int_0^\infty f(\tau) e^{-s\tau}\, d\tau, \qquad (\mathrm{Re}\ s \geq 0). \tag{8.9}$$

The limit can be derived (Franklin and Bair, 1995)

$$\lim_{\omega \to 0} 2\mathrm{Re}\, \frac{\hat{f}(i\omega)}{1 - \hat{f}(i\omega)} = \nu^2\sigma^2 - 1, \tag{8.10}$$

and we will use this value at $\omega = 0$ in equation (8.8) for $S_x(\omega)$. Then we may write

$$S_x(\omega) = |H(\omega)|^2 \nu \left\{ 2\pi\nu\delta(\omega) + \nu^2\sigma^2 + o(1) \right\} \text{ as } \omega \to 0. \tag{8.11}$$

Here $o(1)$ is the generic symbol for some function of ω that tends to zero as $\omega \to 0$; if the lags τ have finite fourth moment $\mathrm{E}(\tau^4)$, it can be shown (Franklin and Bair, 1995) that

$$S_x(\omega) = |H(\omega)|^2 \nu \left\{ 2\pi\nu\delta(\omega) + \nu^2\sigma^2 + O(\omega^2) \right\} \text{ as } \omega \to 0. \tag{8.12}$$

In any case, the lag mean and variance, $\mu = \nu^{-1}$ and σ^2, appear in the power spectral density $S_x(\omega)$ at small frequencies, ω. For the contribution of refraction and other, independent causes, we now recall the formulas

$$\mu = \mu_1 + \mu_2, \qquad \sigma^2 = \sigma_1^2 + \sigma_2^2 \tag{8.13}$$

where the refraction lag τ_{k_1} has mean and variance μ_1 and σ_1^2.

Example 1. To see how these results might be used, suppose that the random refraction lags τ_{k_1} have a gamma distribution with the common p.d.f.

$$f_1(\tau) = \frac{b^a}{\Gamma(a)}\tau^{a-1}e^{-b\tau}, \qquad (\tau \geq 0) \tag{8.14}$$

where a and b are unknown positive constants. Then

$$\int_0^\infty \tau^k f_1(\tau)\, d\tau = a(a+1)\cdots(a+k-1)b^{-k}, \qquad (k = 0, 1, 2, \ldots). \tag{8.15}$$

Then the first and second moments are $\mu_1 = ab^{-1}$ and $a(a+1)b^{-2}$. Since the second moment must equal the variance plus the square of the mean, we find

$$\mu_1 = ab^{-1} \text{ and } \sigma_1^2 = ab^{-2} \tag{8.16}$$

for the common mean and variance of the refraction lags τ_{j_1} $(-\infty < j < \infty)$.

For instance, if μ_1 is fixed and if $\sigma_1 \to +0$, then $f_1(\tau)$ becomes just $\delta(\tau - \mu_1)$.

Suppose the total lag τ_j is the sum of the refraction lag τ_{j_1} and the independent source lag τ_{j_2}, where the source lag has the Poisson p.d.f.

$$f_2(\tau) = \nu_2 e^{-\nu_2 \tau}, \qquad (\tau \geq 0). \tag{8.17}$$

Then the source lags τ_{j_2} have the common mean and variance

$$\mu_2 = \nu_2^{-1} \text{ and } \sigma_2^2 = \nu_2^{-2}. \tag{8.18}$$

Therefore, the total lags τ_j have the common mean and variance

$$\mu = ab^{-1} + \nu_2^{-1} \text{ and } \sigma^2 = ab^{-2} + \nu_2^{-2}. \tag{8.19}$$

The p.d.f. $f(\tau)$ is the convolution $f_1(\tau) * f_2(\tau)$. The Laplace transform $\hat{f}(s)$ is the product

$$\hat{f}(s) = \hat{f}_1(s)\hat{f}_2(s) = \left(\frac{b}{b+s}\right)^a \cdot \frac{\nu_2}{\nu_2 + s}, \qquad (\text{Re } s \geq 0). \tag{8.20}$$

To compute the power spectral density of the neuron output potential $x(t)$, we must use the expression

$$G(\omega) \equiv 2\text{Re}\, \frac{\hat{f}(i\omega)}{1 - \hat{f}(i\omega)}, \qquad (-\infty < \omega < \infty) \tag{8.21}$$

where $G(0)$ is defined as the limit as $\omega \to 0$ (Lemma 1 in Franklin and Bair, 1995):

$$G(0) = \nu^2\sigma^2 - 1 = \mu^{-2}\sigma^2 - 1, \tag{8.22}$$

$$G(0) = (\mu_1 + \mu_2)^{-1}(\sigma_1^2 + \sigma_2^2) - 1. \tag{8.23}$$

The full power spectral density is

$$S_x(\omega) = |H(\omega)|^2 \nu \left[2\pi\nu\delta(\omega) + 1 + G(\omega)\right]. \tag{8.24}$$

For ω near zero, we recall the assertion

$$S_x(\omega) = |H(\omega)|^2 \nu [2\pi\nu\delta(\omega) + \nu^2\sigma^2 + 0(\omega^2)]. \tag{8.25}$$

For instance, if the neuron-firing wave form is just the Dirac δ-function $h(t) = \delta(t)$, then $H(\omega) \equiv 1$, and so

$$S_x(\omega) = \nu[2\pi\nu\delta(\omega) + \nu^2\sigma^2 + 0(\omega^2)] \text{ as } \omega \to 0. \tag{8.26}$$

From experimental data, we can find the power spectrum $S_x(\omega)$. If the graph of

$S_x(\omega)$ near $\omega = 0$ indicates that

$$S_x(\omega) = A\delta(\omega) + C + o(1) \text{ as } \omega \to 0, \tag{8.27}$$

then we may observe the spike strength, A, and the constant C. Knowing A and C, we may write the equations

$$A = 2\pi\nu^2, \qquad C = \nu^3\sigma^2 \tag{8.28}$$

or, equivalently,

$$A = \frac{2\pi}{(\mu_1 + \mu_2)^2}, \qquad C = \frac{\sigma_1^2 + \sigma_2^2}{(\mu_1 + \mu_2)^3}. \tag{8.29}$$

Now, if we know the mean μ_2 and the variance σ_2^2 of the source lags, we can solve the two equations for the mean μ_1 and the variance σ_1^2 of the refraction lags.

By the way, to obtain the lag mean and variance, μ and σ^2, we do not need the power spectral density. If we observe a long sequence of firing instants $t_0 < t_1 < t_2 < \cdots < t_N$, then we may approximate the mean and variance by the empirical mean and variance,

$$\bar{\tau} = \frac{\tau_1 + \cdots + \tau_N}{N} = \frac{t_N - t_0}{N} \tag{8.30}$$

and

$$s^2 = \frac{1}{N-1}\sum_{j=1}^{N}(\tau_j - \bar{\tau})^2 \tag{8.31}$$

where $\tau_j = t_j - t_{j-1}$.

More general models of neuron firing lags. We have assumed that each firing lag τ_j is the sum $\tau_{j_1} + \tau_{j_2}$, where τ_{j_1} comes from refraction and τ_{j_2} comes from independent sources. We have supposed that the variates τ_{j_1}, τ_{j_2} $(-\infty < j < \infty)$ are statistically independent, where each variate τ_{j_1} has p.d.f. $f_1(\tau)$ and each variate τ_{j_2} has p.d.f. $f_2(\tau)$.

More generally, each lag τ_j may need to be considered as the sum of m parts:

$$\tau_j = \tau_{j_1} + \cdots + \tau_{j_m}, \qquad (-\infty < j < \infty). \tag{8.32}$$

Since we still have a stationary renewal process, the lags $\{\tau_j\}$ are still i.i.d., with a common p.d.f. $f(\tau)$ $(\tau \geq 0)$.

Suppose that the partial lag τ_{j_r} has p.d.f. $f_r(\tau)$ $(r = 1, \ldots, m)$. Let the partial lags have expected values and variances

$$\mu_r = \mathrm{E}(\tau_{j_r}) \qquad \text{and} \qquad \sigma_r^2 = \mathrm{E}(\tau_{j_r} - \mu_r)^2, \qquad (r = 1, \ldots, m). \tag{8.33}$$

For $m = 2$ we assumed that the partial lags were independent, so that the total lag τ_j had mean and variance

$$\mu = \mu_1 + \cdots + \mu_m \qquad \text{and} \qquad \sigma^2 = \sigma_1^2 + \cdots + \sigma_m^2. \tag{8.34}$$

Of course, these formulas remain true for $m > 2$ if the partial lags are independent; and the Laplace transform $\hat{f}(s)$ of the p.d.f. $f(\tau)$ is just the product of the Laplace transforms $\hat{f}_r(s)$:

$$\hat{f}(s) = \hat{f}_1(s) \cdots \hat{f}_m(s). \tag{8.35}$$

Even if $m = 2$, these formulas must be changed if the partial lags $\tau_{j_1}, \ldots, \tau_{j_m}$ are dependent. Then we have

$$\mu = \mu_1 + \cdots + \mu_m \qquad \text{but} \qquad \sigma^2 = \sum_{\alpha=1}^{m} \sum_{\beta=1}^{m} \sigma_{\alpha\beta} \tag{8.36}$$

where $\{\sigma_{\alpha\beta}\}$ are the covariances; and now the Laplace transform $\hat{f}(s)$ generally cannot be expressed as the product $\hat{f}_1(s) \cdots \hat{f}_m(s)$.

8.3 Application

We will now apply the results stated in the previous section to neuronal spike trains and will use the equations presented in *Example 1* of the previous section as our model.

8.3.1 Methods

A segment of a spike train is shown across the top of Fig. 8.1. In all analysis presented here, the first 336 msec of the neuronal response during the 2 sec stimulus are discarded to avoid processing the early transient portion of the response. The probability density function for the time between consecutive spikes is estimated from the neuronal data by the interspike interval (ISI) histogram. ISI histograms show the percent of all intervals at each discrete time length (integer values in msec) and are plotted using bars (see Fig. 8.1, left side). The Fourier power spectrum $S_x(\omega)$ is estimated for neuronal spike trains using the methods described in Chapter 7. All power spectra are normalized so that the ordinate value 1.0 corresponds to the mean neuronal firing rate. This normalization is analogous to dividing Eqn. 8.8 by ν. The analytical curves superposed on neuronal power spectra in the figures are described by $\tilde{S}_x(\omega)$, a normalized version of $S_x(\omega)$, which is defined in the Appendix.

8.3.2 Results

The model presented in *Example 1* of the previous section defines each interspike interval to be the sum of a refractory component with a gamma density and an input component with an exponential density. Given the lack of spatial and temporal correlation in the stimulus signal, it is reasonable to first model the neuron input component with the exponential density, consistent with a Poisson process.

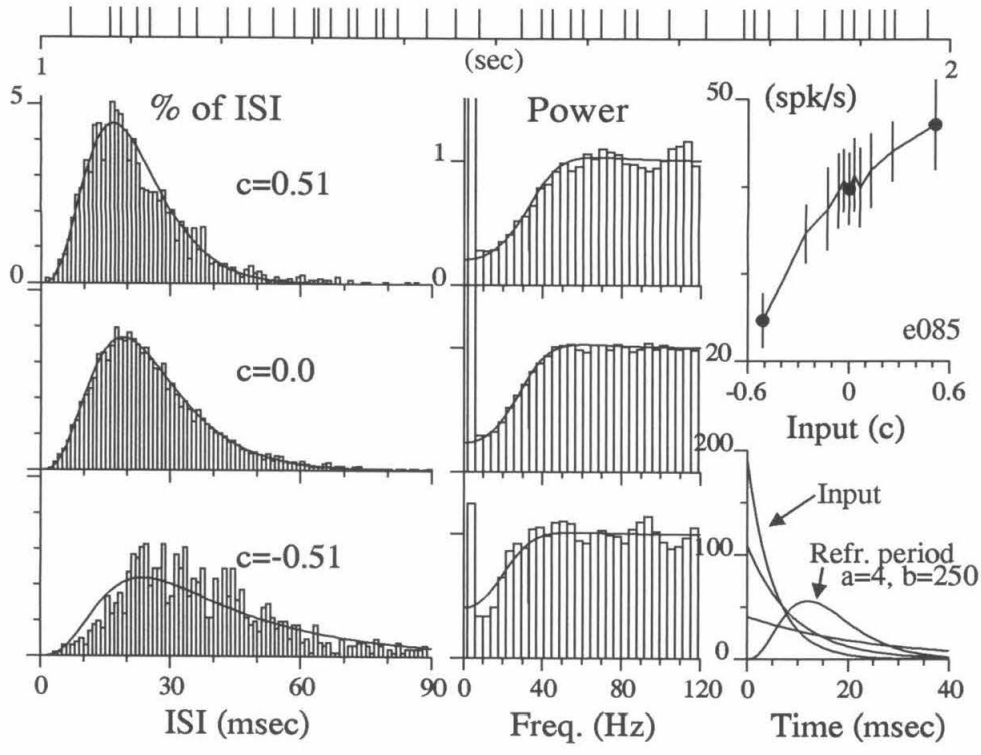

Figure 8.1: Neuronal data vs. theory assuming a fixed refractory period as firing rate changes for cell e085. The top trace shows a one second segment of a spike train recorded at $c = 0$ stimulation. The ISI histogram (left column, bars) and the neuronal power spectrum (center column, bars) were computed from 2 sec duration trials (120 trials at $c = 0$; 30 trials at $c = \pm 0.51$) with the first 336 msec, containing on-transients, discarded. The upper right plot shows the output firing rate of the neuron over the range of input levels tested, and the points mark the cases that are studied here (error bars show standard deviation). Analytical curves for $f(\tau)$ (combine Eqns. 8.5, 8.14, and 8.17) were fitted by eye to the ISI histograms for all 11 input levels plotted in the upper right panel under the constraint that a and b are fixed. For a particular a and b, the parameter ν_2 was set for each input level so that the overall mean firing rate ν matched the values shown for the cell. The lower right plot shows the refraction density, $f_1(\tau)$ (Eqn. 8.14), and input densities, $f_2(\tau)$ (Eqn. 8.17), for the fit. Here $a = 4.0, b = 250.0$, and the values for ν_2 may be read from the plot as the maximum values, at $\tau = 0$, of the exponential densities. Analytical curves for $\tilde{S}_x(\omega)$ (see Eqn. 8.42) based on the parameters from the ISI fits are superposed in the center column. ($h(t)$ is taken to be the δ-function).

For a given neuron, we determine the values of the refractory period parameters a and b which give the best fit by eye to the series of ISI histograms at all c values available. Note that a and b do not vary as a function of the input c—only ν_2 is allowed to vary but is constrained so that

$$\mu_{est}(c) = ab^{-1} + \nu_2^{-1}(c), \tag{8.37}$$

where μ_{est} is the measured mean spike rate as a function of input strength c (Fig. 8.1, upper right) and the notation $\nu_2(c)$ indicates the input dependence of ν_2. Fig. 8.1 shows the ISI histograms and their fits $f(\tau)$ (left column) for $a = 4.0$, $b = 250.0$ (where a/b is the mean ISI in seconds and a/b^2 is the variance) at the three c values indicated by dots in the upper right plot. The lower right panel of Fig. 8.1 shows the refractory period density, $f_1(\tau)$, and the three input densities, $f_2(\tau)$, for the fit parameters. Roughly half of 58 cells examined showed fits qualitatively similar to that shown here.

The center column of Fig. 8.1 shows the measured power spectra of the neuronal spike trains (bars) and $\tilde{S}_x(\omega)$ (smooth curve) based on the parameters a, b, $\nu_2(c)$ (see Eqn. 8.42 in the Appendix for an explicit formula for $\tilde{S}_x(\omega)$). The refractory period causes a dip in the power spectra at low frequencies—without the refractory period, the spectrum, that of Poisson impulses, would be flat.

Fig. 8.2 shows the results of applying a similar analysis to a neuron that has a shorter refractory period ($a/b = 4.7$ msec rather than 16 msec) and a wider range of output firing rates (upper right panel). It is striking that the left (rising) sides of the ISI histograms are well fit by the same fixed refractory period ($a = 8$, $b = 1700$, thus $\mu_1 = 4.7$ msec, $\sigma_1 = 1.7$ $msec$) over a broad range of spike rates, particularly for $c = 1.0$ where the firing rate appears, based on the upper right plot (Fig. 8.2), to be saturated. The analytical power spectra (smooth curves on center column plots)

predict the shape change observed in the neuronal power spectra (bar plots, center column Fig. 8.2) as c changes. Again, the values of the theoretical inputs, ν_2, are observed as the y-intercept of the decaying exponentials in the lower right plot. Some small systematic errors appear in the fits, in particular, the right side of the ISI histograms tends to have a faster than exponential fall off.

Some cells tend to fire bursts of action potentials (see Chapter 7), as shown in the top spike train of Fig. 8.3, which may violate the assumption that the spike train $x(t)$ is a renewal process. The theory can still be applied by considering bursts to be "events" and operating on inter-*event* intervals (IEIs) rather than ISIs. As in Chapter 7, an event is defined as the longest train of consecutive action potentials that have no ISIs greater than d msec ($d = 8$ msec here). We define the process $y(t)$, composed of δ-functions at the center of each event, to be the event train corresponding to the spike train $x(t)$. The event train is plotted beneath the spike train at the top of Fig. 8.3. The spike train power spectrum (Fig. 8.3, center) has a substantial peak near 40 Hz which is not present in the spectrum for events. The IEI histogram (lower left, bars) was fitted to the same mathematical model used in the previous figures, and the analytical power spectra $\tilde{S}_y(\omega)$ (smooth curves) from this model match closely those estimated from the event trains. To derive an analytical spectrum for the spike trains, we define $h(t)$ to be the boxcar function (see inset below the event train in Fig. 8.3) which serves as a model for the prototypic event. The analytical spectrum, $|H(\omega)|^2 \tilde{S}_y(\omega)$, using a deterministic approximation to the variable event shape, provides an excellent prediction to the location and shape of the peak in the power spectrum. This example shows that the presence of bursts and an inter-burst refractory period can combine to produce a maximum in the power spectrum.

Fig. 8.4 demonstrates a case in which a peak in the power spectrum develops from regular firing rather than from burst firing. Again we reduce $x(t)$ to the event trains $y(t)$, fit the model to the IEI histograms (left column), and plot the derived power

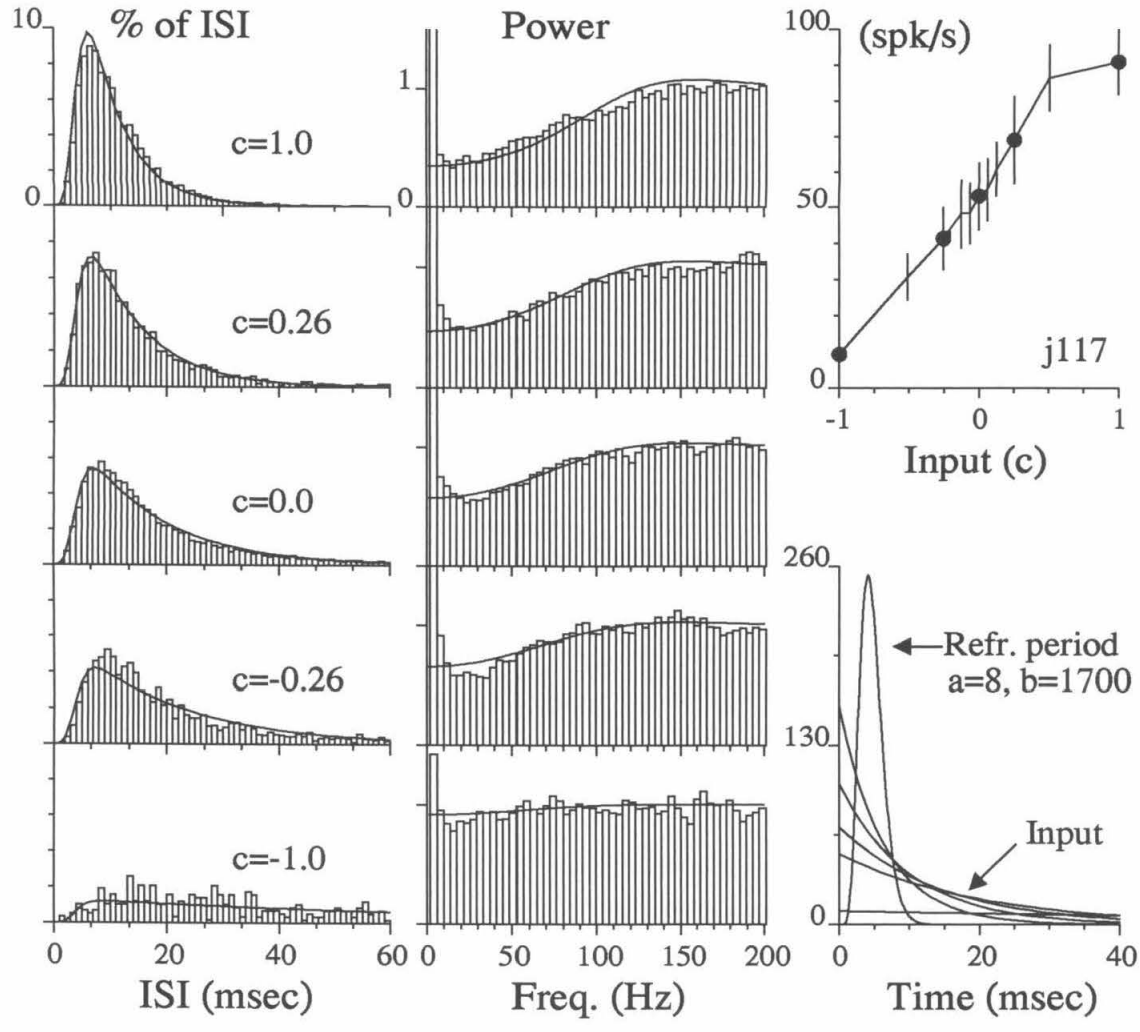

Figure 8.2: Neuronal data vs. theory for a neuron (j117) with a short refractory period and a broad range of output firing rates. The arrangement is similar to that in Fig. 8.1. Under the assumption of a fixed refractory period ($a = 8.0$, $b = 1700.0$), the left sides of the fits to the ISI histograms (left column) match the data well as firing rate changes (upper right) and even after the firing rate appears to saturate at $c = 1.0$. The analytical curves for $\tilde{S}_x(\omega)$ (solid lines, center column) give good approximations to the shape of the neuronal power spectrum. (Forty trials of data were used at each stimulus level, except 80 trials were used at $c = 0$.)

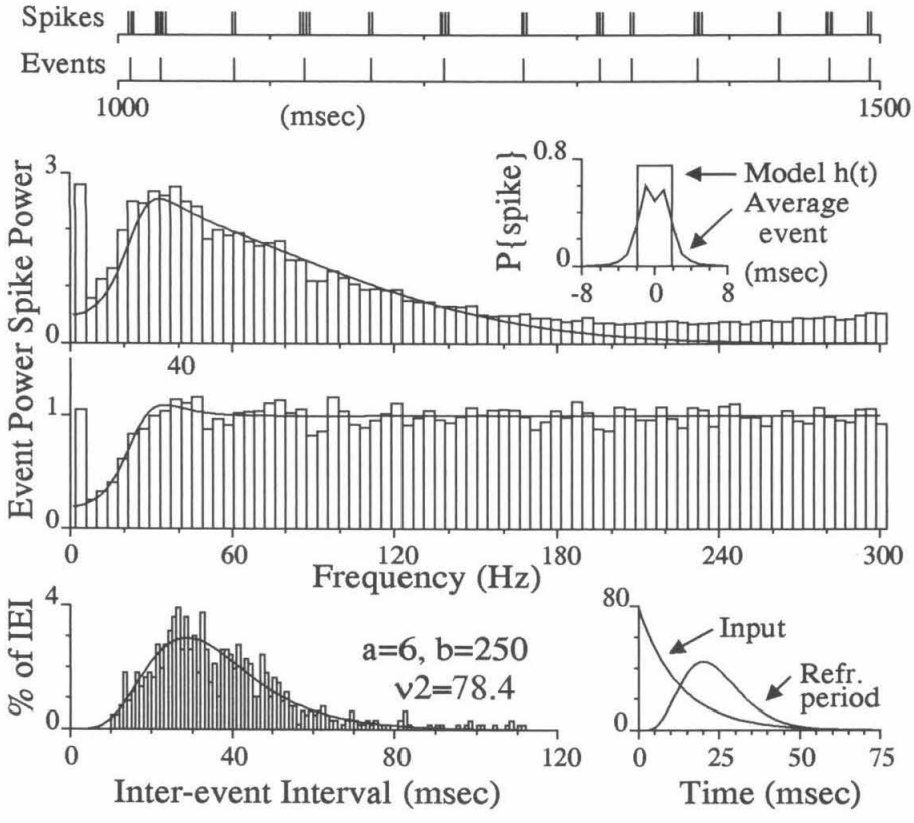

Figure 8.3: Modeling bursts of action potentials as single "events" allows the theory to predict the peak in the power spectrum which occurs for cells that fire bursts. The top trace shows a spike train fired by neuron j001. Directly below it is the event train, a series of δ-functions at the center of each event, where an event is defined as the longest train of consecutive action potentials with no ISIs greater than d msec ($d = 8$ here). The power spectrum of the spike trains at $c = 0.51$ has a peak near 40 Hz (upper spectrum) while the power spectrum of the corresponding event trains is flat at frequencies higher than 40 Hz (lower spectrum). This indicates that the large peak observed in the power spectrum is due to the shape, not the timing, of the events. A model of the inter-event interval (IEI) histogram (lower left) based on a refractory lag and an input lag (lower right) leads to an analytical prediction $\tilde{S}_y(\omega)$ (smooth curve) of the event power spectrum. When events are modeled by setting $h(t)$ to the boxcar function shown in the inset below the event train, the analytical expression $\tilde{S}_x(\omega)$ for the spike power spectrum predicts the peak observed in the data (upper spectrum). Note, Eqn. 8.42 in the Appendix must be modified to account for the boxcar $h(t)$.

spectra $\tilde{S}_y(\omega)$ against the neuronal spectra (right column). The peak in the event power spectrum is observed to grow as the event firing rate ν_2 increases. The values of ν_2 for $c = 0.0,\ 0.51,\ 1.0$ are approximately 25, 50, and 100 events/sec, respectively (see y-intercepts of the exponential densities in the inset, Fig. 8.4). In the limit as $\nu_2 \to \infty$, $f(\tau) \to f_1(\tau)$, i.e. the interval p.d.f. becomes the refractory period p.d.f. If $f_2(\tau)$ is regular enough to produce a peak in the power spectrum in this limiting case, then we may observe peaks as seen in the right column of Fig. 8.4. However, if $f_2(\tau)$ is, for example, an exponential distribution ($a = 1$), the power spectrum will remain flat even as $\nu_2 \to \infty$.

A comparison of the refractory period p.d.f.'s for the four neurons analyzed in the previous figures is shown in Fig. 8.5. The means, a/b, of the gamma densities range from 4.7–24 msec. Although the means for cell e039 and e085 are 12 and 16 msec, their standard deviations are 3.5 and 8.0 msec, respectively, indicating the usefulness of having two parameters for the refractory p.d.f. The range of values here are typical of those for the 58 cells analyzed.

8.3.3 Discussion

Fig. 8.1 and Fig. 8.2 demonstrate that a fixed stochastic refractory period can be adequate to account for the distribution of interspike intervals across a wide range of firing rates when the input is assumed to induce the neuron to fire at Poisson time instants. (However, Turcott et al., 1994, report that the refractory period can change over time in auditory neurons.) The power spectra of the neuronal data in Fig. 8.1 and Fig. 8.2 are consistent with the analytical curves from the formula derived here, Eqn. 8.8, for the power spectrum of a renewal process with a specified interval density, $f(\tau)$. In particular, the presence of a refractory period causes a dip in the power spectrum at low frequencies, consistent with the results of Chapter 7.

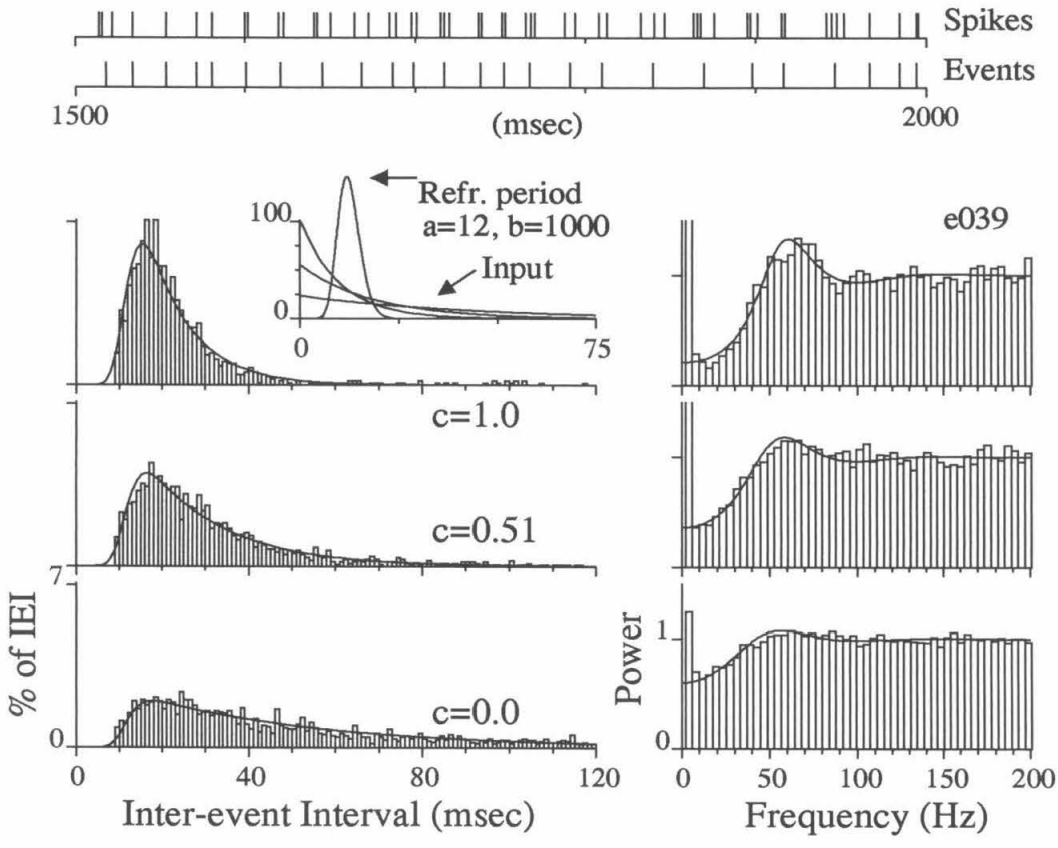

Figure 8.4: Spectral peaks due to regular firing are different from those caused by bursts and may occur for a renewal process as the input rate becomes fast with respect to the refractory period. The methods of Fig. 8.1 and Fig. 8.3 are combined to model the IEI and event power spectrum for neuron e039. The power spectra (bars, right column) are computed for event trains, and the analytical predictions $\tilde{S}_y(\omega)$ are based on the δ-function model of events to eliminate the peak due to bursts. As the event rate increases for stronger stimuli, a peak, predicted by the model, develops in the power spectra near the inverse of the mean IEI. Peaks in the power spectra such as these result from the regularity of the timing of the events, rather than from the event shape as seen in Fig. 8.3.

Fig. 8.3 and Fig. 8.4 show two ways for peaks to arise in the power spectrum of spike trains. First, a peak may arise due to the compound effect of a dip at low frequencies created by a refractory period and the attenuation at high frequencies induced by firing bursts of action potentials (Fig. 8.3). The frequency of the peak is a function of the refractory period parameters and the shape of the bursts, i.e. events, and is not related to regularity in the inter-event interval. The second type of peak is caused by regularity in the ISI or IEI density, and this regularity is well modeled as the result of an increasing firing rate with a fixed refractory period density. Fig. 8.4 demonstrates this for inter-event intervals of a neuron that fires bursts; however, we have observed the same effect in neurons that fire isolated action potentials.

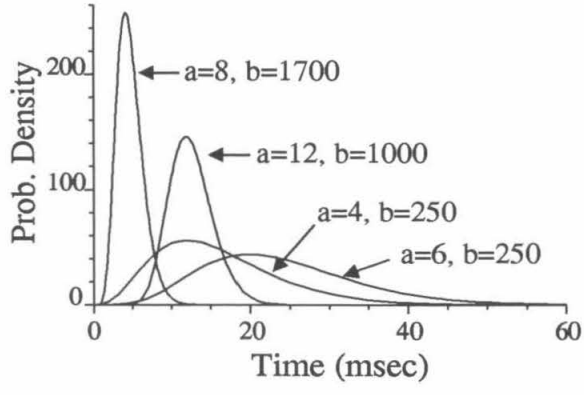

Figure 8.5: A comparison of the refractory period densities, $f_1(\tau)$, for the four cells studied in the previous figures. Although the mean μ_1 and the standard deviation σ_1 of the refractory period often rise or fall together, the middle two examples here have somewhat similar means, 12 and 16 msec, but quite different standard deviations, 3.5 and 8 msec. Recall $\mu_1 = a/b$, $\sigma_1^2 = a/b^2$.

We have presented two mathematical models of spike trains which can lead to peaks near 40 Hz in the power spectrum based on the formulas derived here and which both incorporate the notion of a refractory period. We emphasize that the concept of refractory period used here is not restricted to a neuron's intrinsic limitations for

quickly firing a second action potential; we do not know the pre-synaptic input to these neurons and cannot distinguish intrinsic from network or input effects which may cause a paucity of short ISIs.

8.4 Appendix

The spike-rate normalized power spectrum $\tilde{S}_x(\omega)$ is derived from formula (8.8) by setting $h(t) = \delta(t)$, dividing by the mean spike rate, ν, and subtracting the δ-function term. Thus,

$$\tilde{S}_x(\omega) = 1 + 2\mathrm{Re}\,\frac{\hat{f}(i\omega)}{1-\hat{f}(i\omega)}, \qquad (-\infty < \omega < \infty,\ \omega \neq 0), \tag{8.38}$$

where $\tilde{S}_x(0) = \nu^2\sigma^2$, as defined by continuity in equation (8.10). Let the interval density $f(\tau)$ be defined as the convolution of a gamma density, given in (8.14), and an exponential density, given in (8.17) so that its Laplace transform $\hat{f}(s)$ is, as given in formula (8.20),

$$\hat{f}(s) = \left(\frac{b}{b+s}\right)^a \cdot \frac{\nu_2}{\nu_2+s}, \qquad (\mathrm{Re}\ s \geq 0). \tag{8.39}$$

Then, $\hat{f}(i\omega)^{-1} = \rho e^{i\phi}$, where

$$\phi = a\tan^{-1}(\omega/b) + \tan^{-1}(\omega/\nu_2), \tag{8.40}$$

and

$$\rho = \left(\frac{b^2+\omega^2}{b^2}\right)^{a/2} \sqrt{\frac{\nu_2^2+\omega^2}{\nu_2^2}}. \tag{8.41}$$

Under the previous assumptions and definitions, an explicit expression for $\tilde{S}_x(\omega)$ is

$$\tilde{S}_x(\omega) = 1 + 2\frac{\rho\cos\phi - 1}{\rho^2 - 2\rho\cos\phi + 1}, \qquad (-\infty < \omega < \infty,\ \omega \neq 0). \tag{8.42}$$

Equation (8.42) depends on only a, b, and ν_2 and is the expression for the analytical curves superposed on the neuronal power spectrum histograms in the figures in this paper (except for the top spectrum in Fig. 8.3, which uses a different definition for $h(t)$).

Chapter 9

The Local Field Potential

This chapter summarizes analysis of a small database of local field potential (LFP) signals that were recorded in areas MT and MST of a behaving macaque monkey. Experiments were conducted by Gregory Horwitz in Prof. William Newsome's Laboratory at the Stanford University School of Medicine. As mentioned in Chapter 1, our initial motivation for studying the temporal structure of spike trains was to examine 40 Hz oscillation. Kreiter and Singer (1992) reported that oscillation was present in area MT and MST of the awake macaque monkey, but was more variable in frequency and duration than the 40 Hz oscillation in cat. We found no striking evidence for oscillation in our analysis of single unit data in Chapter 7, but because oscillations are more frequently observed in the LFP than in single spike trains (e.g., Eckhorn et al., 1988), we wanted to examine LFP data.

9.1 Background and Summary

Local field potentials are signals recorded between 1–100 Hz from a single extracellular electrode that presumably reflect the integration of membrane currents in a local region of cortex (e.g., Mitzdorf, 1985). In the anesthetized cat, LFP receptive fields

are usually smaller than 1.2 mm of cortex in diameter, and it has been argued that oscillation in the LFP is of cortical origin and may result from superposition of local excitatory dendritic potentials (Eckhorn et al., 1988; Gray and Singer, 1989). It is not uncommon that the stimulus specificity of the LFP is different than that of single units recorded from the same electrode (Gray and Singer, 1989). A recent report demonstrated that the LFP is orientation-tuned with 2° receptive fields at most parafoveal sites in anesthetized macaque V1 (Victor et al, 1994). However, Mitzdorf et al. (1994) reported that diffuse luminance changes are more effective than oriented bars in driving LFPs in rabbit visual cortex. Oscillatory LFPs in V1 shown in the literature are commonly on the order of 100 μV peak-to-peak.

This analysis covers three aspects of the data: (1) the relationship between the average LFP and the post-stimulus time histogram (PSTH) for an isolated neuron, (2) the relationship between the LFP and the spike occurrence times within a single trial, including spike-triggered averages of the LFP, and (3) the power spectrum of the LFP.

We found that the LFP had stimulus-locked fluctuations near 50 Hz in the first 100 msec of the response. Spike-triggered averaging revealed a characteristic bi-phasic LFP waveform lasting about 10 msec that was associated with the occurrence of an action potential. The amplitude of this waveform varied greatly across sites recorded here. A prediction of the LFP based on the PSTH and the spike-related LFP waveform did not account well for the actual LFP. The power spectra of the LFP were greatly contaminated by electrical interference at 60 and 180 Hz, and there was an undesirable attenuation of low frequencies due to filters in the recording setup. Because of the small database examined here and the 60 Hz interference in the recording setup, we could not, unfortunately, draw conclusions about the presence of 40 Hz oscillation in the LFP.

9.2 Methods

The LFP was recorded at 5 sites (2 MT, 3 MST) in a behaving macaque monkey. The visual stimulus was a dynamic random dot display as described in Chapter 2, except that the stimulus duration was typically limited to 1 sec. Electrodes were Parylene insulated tungsten with 1 kHz impedances typically in the range of 0.5–2.0 MΩ. The LFP signal was sampled at 500 Hz. The absolute scale of the LFP plots in this report is unknown—there were a number of variable gain amplifiers between the electrode and the A/D converter which sampled the data.

A low pass filter was added to the recording setup to prevent aliasing in the digitized LFP signal. A schematic of the setup is shown in Fig. 9.1. A characterization of the low pass filter is presented in Fig. 9.2. The filter causes a delay of 5–6 msec between the spike train recordings and the LFP recording; therefore, the LFP plots, for example in Fig. 9.3, are shifted to the left by 6 msec.

9.3 Stimulus-Locked LFP Fluctuations

Figs. 9.3–9.5 show examples of PSTHs and average LFPs for varying levels of motion coherence at four recording sites. The fifth recording site was not included in this discussion because the LFP was dominated by the action potentials of a single cell (see the top of Fig. 9.8, and the section on spike-triggered averaging). Histograms are used to represent the PSTH for $c = \pm 1.0$, and the corresponding LFP plots are superimposed below the PSTHs. The thick bars at $t = 0$ on the LFP indicate a typical standard deviation above and below the mean. In Fig. 9.3, stimulus-locked modulation of the LFP begins at $t = 55$ msec.[1] The time course of the modulation

[1]Greg Horwitz has estimated that the stimulus onset is biased by 25 msec; therefore, the LFP response occurs with a latency of 30 msec. The 25 msec offset is **not** taken into account in this report. Maunsell and Gibson (1992) report that the shortest latencies to striate cortex were 20 msec

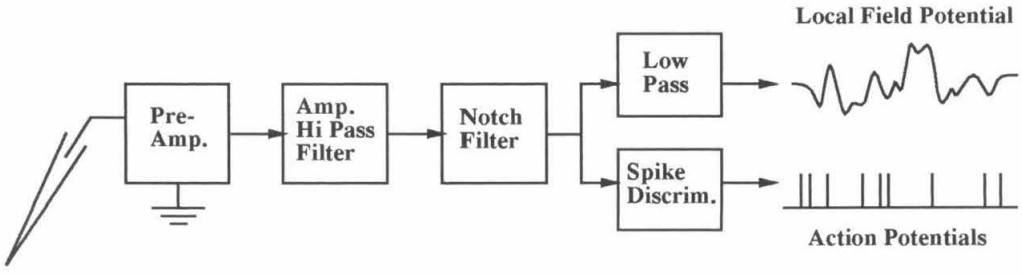

Figure 9.1: The recording setup. The signal from the electrode (left) was first amplified by the head-stage pre-amplifier. Further amplification and high pass filtering occurred and a notch filter was used to eliminate the eye-coil signal. Action potentials were discriminated from this signal using a time-amplitude window; only the occurrence times were recorded. A low pass filter (for a characterization, see Fig. 9.2) further processed the signal before the LFP was digitized. Aluminum foil shielding was placed around the head-stage pre-amplifier to reduce the amplitude of 60 Hz noise.

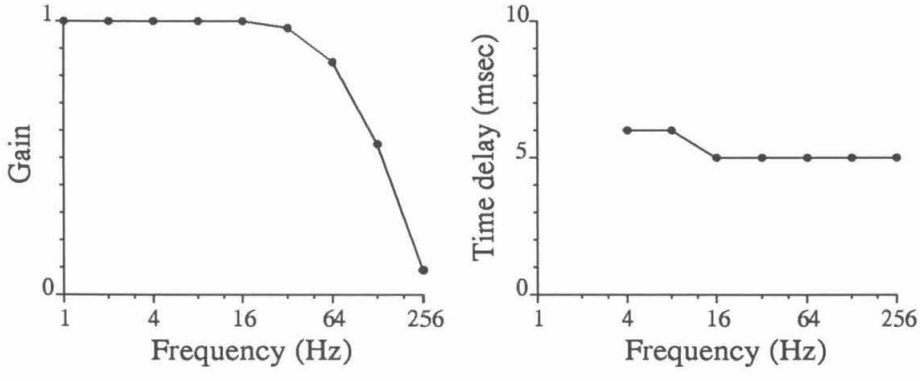

Figure 9.2: Characterization of the low pass LFP filter. This filter was between the notch filter and the LFP recorder. A 200 mV peak-to-peak sine wave was used to characterize the gain and time delay of the filter. The low pass filter box was set to 100 Hz. At this setting, the LFP was sampled at 500 Hz, since anything below half that frequency was highly attenuated (left). The filter introduced an approximately constant 5–6 msec time delay over the frequency range of interest. This filter was responsible for the high-frequency attenuation observed in Fig. 9.11.

is similar for motion in either direction for approximately 100 msec, and therefore is not directional.

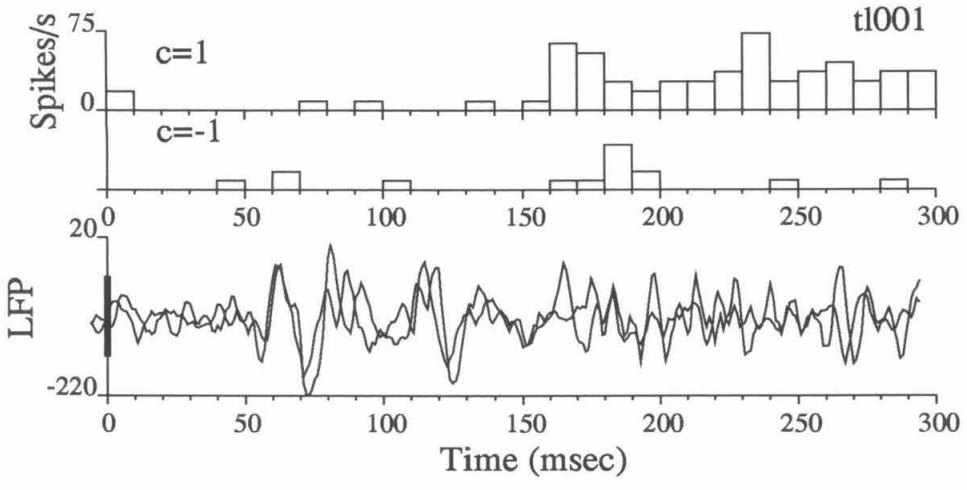

Figure 9.3: PSTHs (upper plots) for cell tl001 and accompanying LFPs (bottom) for coherent motion in the preferred and null directions at site tl001 in area MT. The cell's response begins at 160 msec and is directional (see PSTHs). Earlier evidence of the stimulus appears in the average LFPs which show large amplitude fluctuations starting at 55 msec. These fluctuations appear non-directional from 55–160 msec. After 160 msec, the LFP shows larger amplitude fluctuations for $c = 1$ than for $c = -1$, and is in this sense directional. The thick bar at the left on the lower plot indicates a typical standard deviation above and below the mean LFP ($n = 11$ for $c = 1$, $n = 12$ for $c = -1$).

A detailed description of each site will not be given, but a few general observations for sites tl001–4 can be made. (1) The LFP has a non-directional, stimulus-locked pattern of modulation during the period from approximately 50 to 150 msec. (2) The modulation of the LFP during this period shows no strong resemblance to the single cell PSTHs, which are also non-directional during this time. (3) There is some consistency in the pattern of the modulation across sites, and this is demonstrated in Fig. 9.6 which shows the LFP averaged across all c values for each recording site.

in one animal (M. fascicularis) and 30 msec in another (M. mulatta).

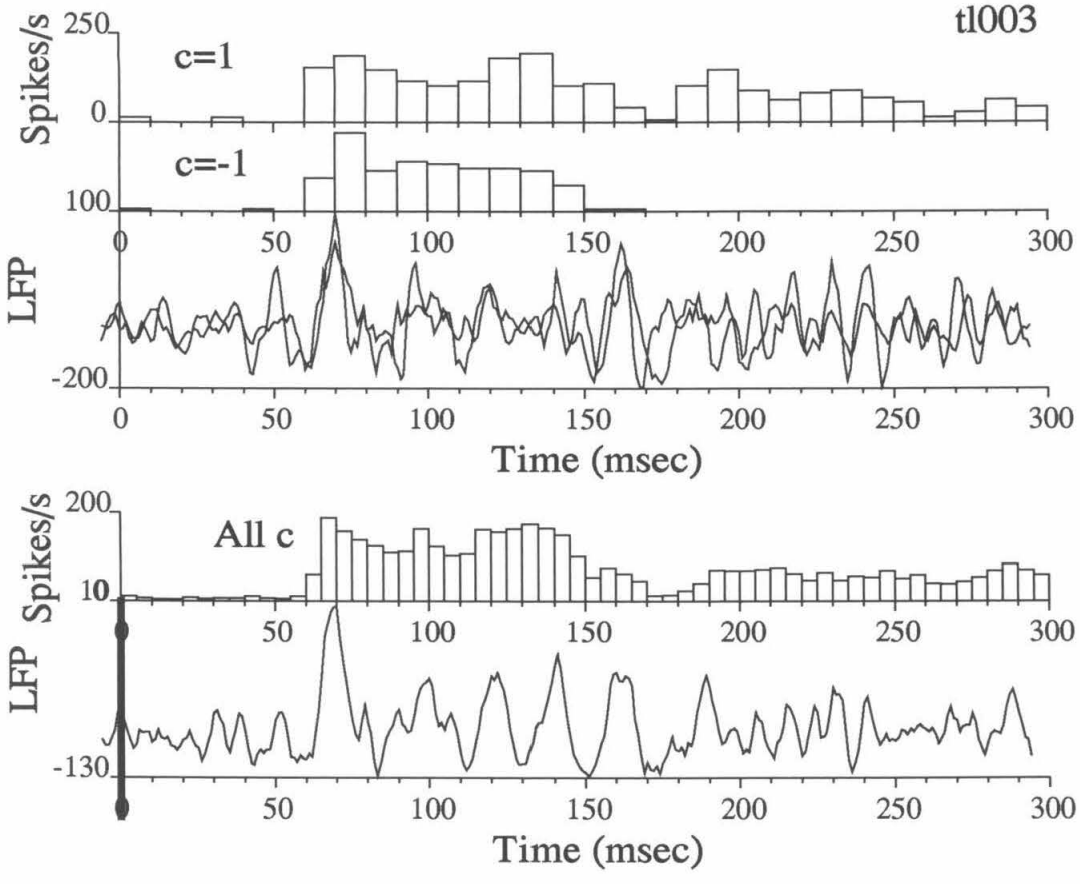

Figure 9.4: Site tl003, area MST. The top is similar to Fig. 9.3. The bottom two plots show the average for all coherence values. The PSTHs and LFPs show a non-directional response from 60–150 msec. The averages for all c show that the oscillation in the LFP is not present in the PSTH for the cell. (Ten trials for each of 8 coherence levels.)

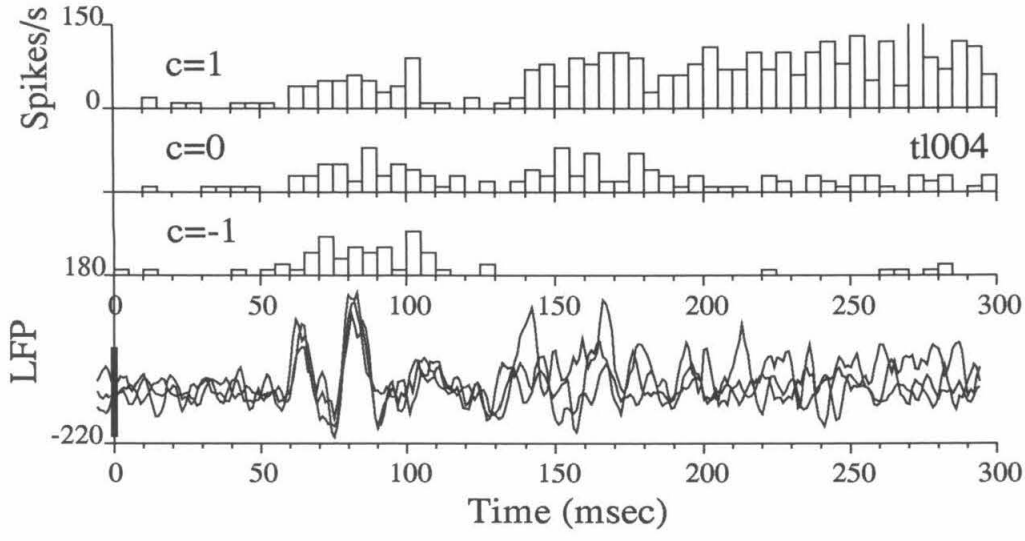

Figure 9.5: Site tl004, area MT. Three coherence levels ($c = 1, 0, -1$) are used here. Responses are non-directional from 50–135 msec. The LFP fluctuations are not present in the PSTH. (Twenty trials at each coherence level.)

In Fig. 9.6, the initial LFP peaks for sites tl001, tl002, and tl004 occur at similar times (top), while the initial peak for tl003 occurs somewhat later (bottom). There is a typical 20 msec spacing between LFP peaks, which corresponds to 50 Hz. This is apparent from the peaks at 60 and 80 msec (top) and 100, 120, 140, and 160 msec (bottom).

Is there stimulus-locked modulation in the LFP after the initial non-directional period? For site tl003, a single neuron was well isolated and repeated trials were performed with the same stimulus seed. Stimulus-locked modulation is present in the PSTHs in Fig. 9.7,[2] but the noise level is too high to conclude whether the LFP is related to the PSTH during the sustained period of the response.

[2]Notice the preferred and null direction stimuli yield similar time courses, as in Fig 4.1.

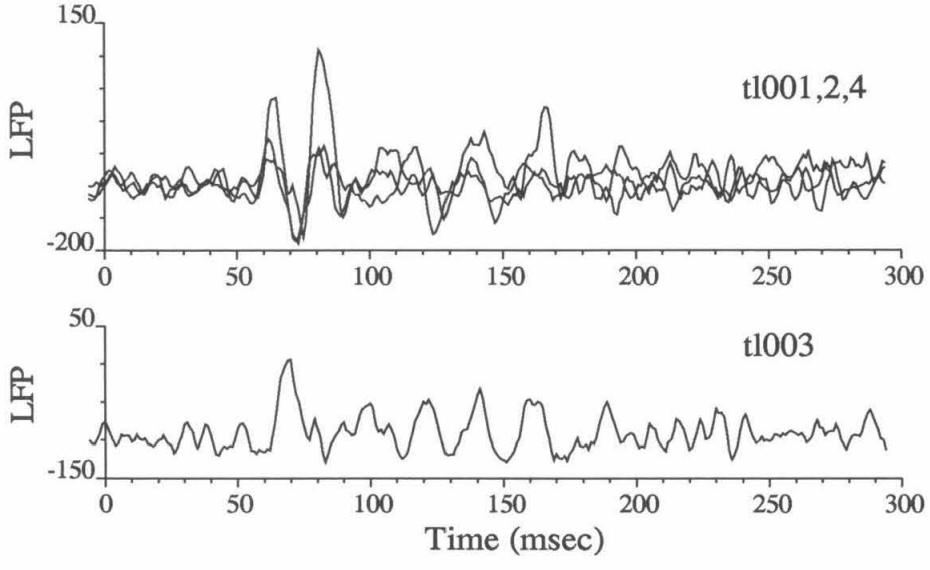

Figure 9.6: The LFP averaged across all c at 4 sites. The top shows sites 1, 2, and 4, which had similar time courses from 60–90 msec. The lower plot shows (using the same vertical scale) the LFP for site 3 which had its initial peak near 70 msec. It is interesting that the initial peaks near 60 and 80 msec in the upper plot and the later peaks in the lower plot at 100, 120, 140, and 160 msec all had approximately a 20 msec separation, corresponding to 50 Hz. Site 5 was not included since the LFP was merely a reflection of the PSTH for the single unit (see Fig. 9.8).

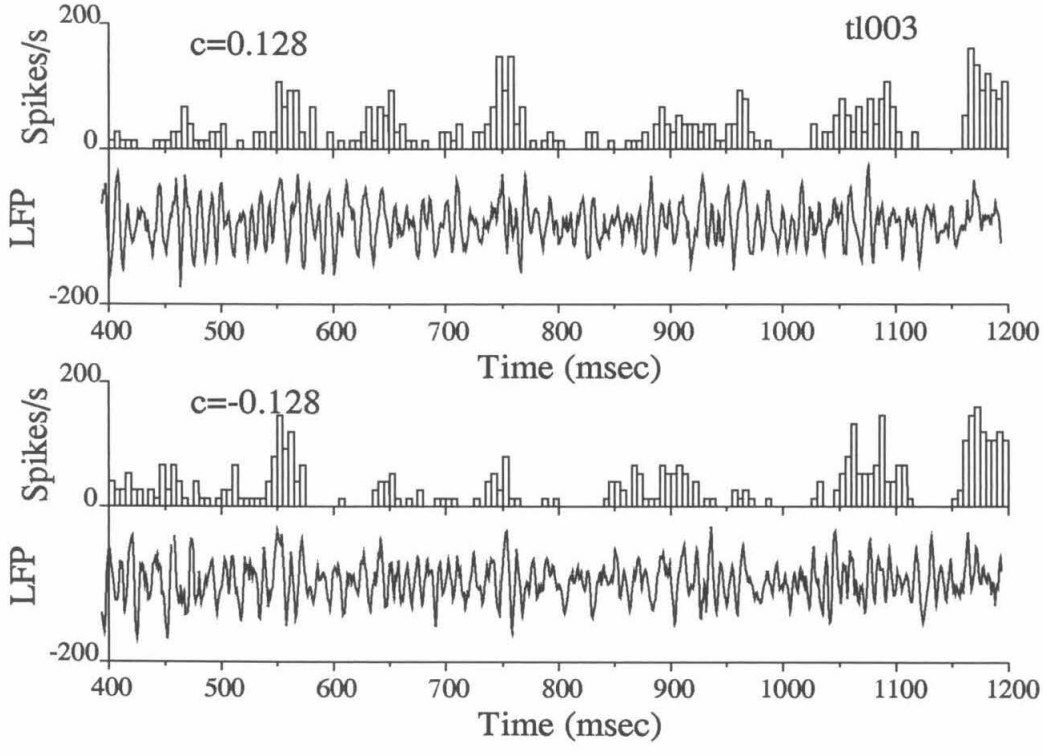

Figure 9.7: PSTH and LFP for repeated identical stimulation of tl003. The upper two plots show the PSTH and average LFP for $c = 0.128$, while the lower two plots show data for $c = -0.128$. The PSTHs (bars) are somewhat similar for the preferred and null direction of motion because only 13% of the stimulus dots, the signal dots, are different. There is no strong relationship between the amplitude of modulation of the LFP and the peaks in the PSTH. The visual stimulus ends at 1000 msec.

9.4 Spike-Triggered Average LFP

For recording site tl005, the LFP is dominated by the action potentials of a single neuron. When the single unit spike train is aligned with the LFP recorded on the same trial, as in Fig. 9.8 (top), a one-to-one correspondence is noted between spikes and large peaks in the LFP. This was not the case with the other four recording sites, and one example of this is shown for comparison at the bottom of Fig. 9.8. For each site, the spike-triggered average LFP was computed to show explicitly the average modulation in the LFP associated with an individual action potential. Our spike-triggered averages, shown in Fig. 9.9, had a peak centered at the spike time (once the 5–6 msec filter delay is subtracted) and a negative-going lobe of smaller amplitude following the spike. This is similar to data shown in Fig. 6A of Eckhorn et al. (1988).

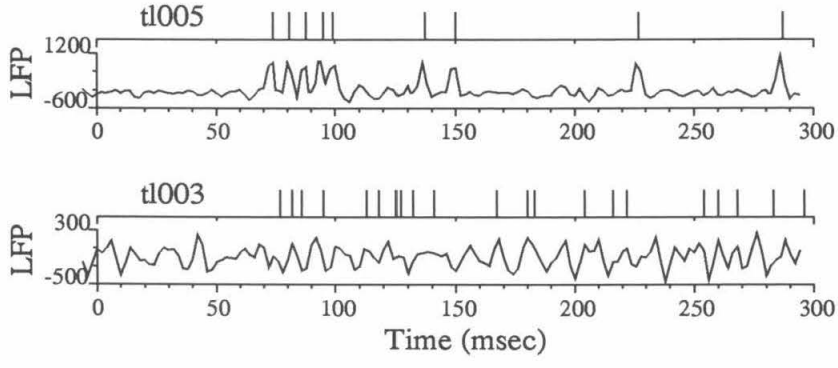

Figure 9.8: For single trials at two sites, the spike trains are plotted above the LFPs. Top: site tl005, area MST. The LFP for tl005 has prominent peaks that align with each action potential recorded for the single unit. The average LFP and the PSTH therefore look very similar and are not shown. Part of a single trial from site tl003 is shown (bottom), for which there is no obvious correspondence between peaks in the LFP and action potentials.

Using the time of occurrence of each action potential as the origin, the average LFP is computed for 5 sites and plotted in Fig. 9.9. For site tl005, where each action *potential* appears unmistakably in the LFP, the spike-triggered average LFP is the

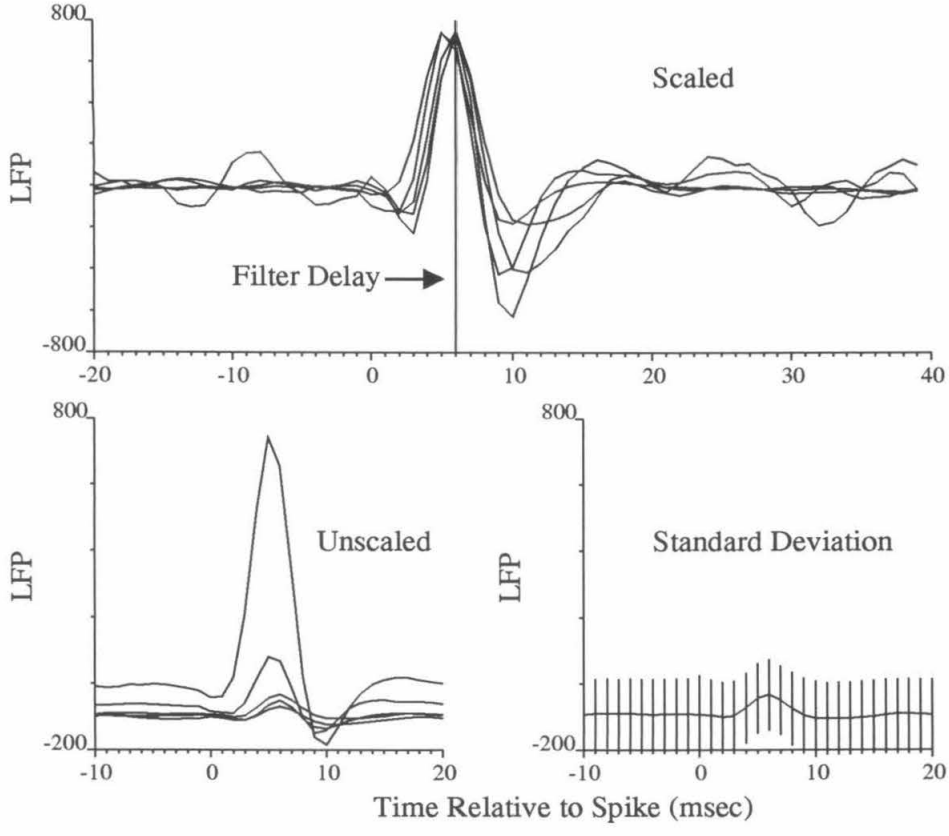

Figure 9.9: The average LFP is shown triggered on action potentials of single units for the 5 recording sites. The spike-triggered LFP averages were scaled to the have the same peak value (top) and unscaled (bottom left). The error bars in the lower right panel show one standard deviation above and below the mean for tl003. The tallest peak (lower left) corresponds to tl005, while the second tallest is from tl004 (recorded with the same electrode, believed to be of lower impedance than the electrode used at the other 3 sites). The low pass filter between the notch filter and LFP recorder caused a 5–6 msec delay between the spike time, $t = 0$, and the recorded peak in the LFP.

tallest waveform in the bottom left of Fig. 9.9. However, at the other four sites where action potentials are not apparent in the LFP (as seen in Fig. 9.8, bottom), the spike-triggered waveform is smaller. In the top of Fig. 9.9, the spike-triggered averages are scaled to have the same peak height, demonstrating that the ratio of the peak to the subsequent dip is variable across sites. At the bottom right of the same figure, the mean waveform for tl003 is shown with error bars indicating plus and minus one standard deviation.

The waveforms of action potentials recorded extracellularly are typically less than 1 msec in duration, while these spike-triggered averages are on the order of 10 msec. Also, the 5–6 msec delay of the low pass filter is noted on the upper panel of Fig. 9.9. Taking this into account, the peaks of the spike-triggered LFP waveforms align with the peak of the recorded action potential. We are uncertain of the origin of the shape of the spike-triggered averages.

9.5 LFP Prediction from Spike Trains

Using the spike-triggered average LFP waveforms from the previous section, a prediction can be made of the average LFP based on the PSTH (binned at the original millisecond sampling resolution) at a given coherence level. Data from site tl003 is used because the single unit isolation was excellent and repeated trials were performed with identical stimuli. The predicted LFP waveform is computed in essence by replacing each action potential in the spike trains with a model of the spike-triggered LFP waveform (see Fig. 9.10, top) and averaging over all spike trains resulting from a particular stimulus. Fig. 9.10 (bottom) shows a segment of the predicted LFP (thick line) superimposed on the actual LFP. The predicted LFP over 50–170 msec is suggestive of the modulation in the actual LFP, but severely underestimates the modulation amplitude. This indicates that the LFP signal is not likely to be associated with only

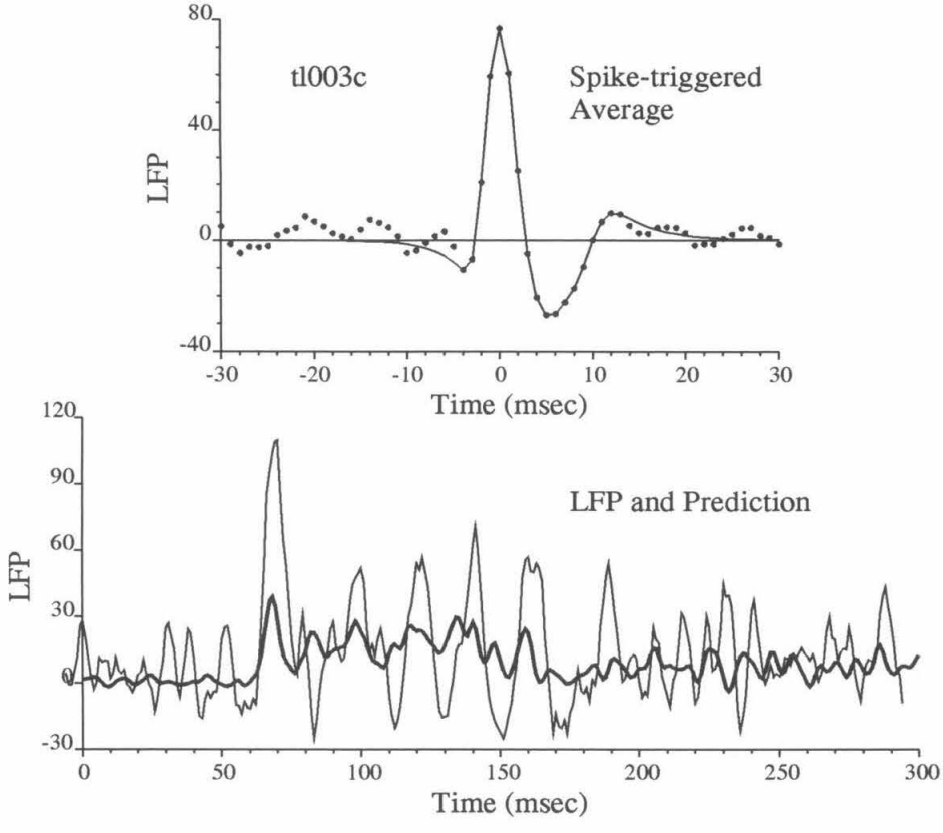

Figure 9.10: Attempted prediction of the LFP from the PSTH using the spike-triggered average LFP. The line in the upper plot is a model of the contribution of each action potential to the LFP. The dots show the spike-triggered average LFP computed for the sustained period of tl003's response to visual stimulus (averaged across motion coherence levels). The lower plot shows the average LFP across all coherence levels (thin line) and the prediction (thick line) based on the convolution of the PSTH and the model kernel in the upper plot. The prediction greatly underestimates the amplitude of the oscillations in the early portion of the response. This is consistent with the notion that stimulus-locked signals other than those directly related to the occurrence of an action potential in tl003 are contributing to the LFP signal.

the action potentials of this single unit, as is consistent with the mismatch in shape between the PSTHs and LFPs in Fig. 9.4. Indeed, the majority of the LFP is likely to be dendritic in origin, given active dendritic currents and the fact that up to 98% of the neuronal membrane is in the dendrites.

9.6 LFP Power Spectra

Fig. 9.11 shows power spectra of the LFP for two recording sites. The two most striking features of these spectra are peaks due to electrical interference and the envelope created by analog filters (see Fig 9.1). The most prominent peaks occur at 60 and 180 Hz. The envelope of the spectrum is a result of a high pass filter that precedes action potential discrimination and a subsequent low pass filter recently added for LFP recording. (Fig. 9.2 shows a characterization of the low pass filter made using at 200 mV peak to peak sine wave.)

The top of Fig. 9.11 shows average spectra at site tl003 for 8 coherence levels, $c = \pm 1.0$, ± 0.51, ± 0.26, and ± 0.13. The lower part of the figure shows spectra for $c = 0$ and for $c = \pm 1.0$ at site tl004. For tl004 there appears to be a relationship between the motion coherence, thus the firing rate of the cell, and the integrated power. Note that tl004 was stimulated with a fixation paradigm that had a 2 sec motion stimulus and no delay period.

The electrical interference and the undesirable low frequency attenuation due to filtering make the data unsuitable for drawing conclusions about the presence of 40 Hz oscillation.

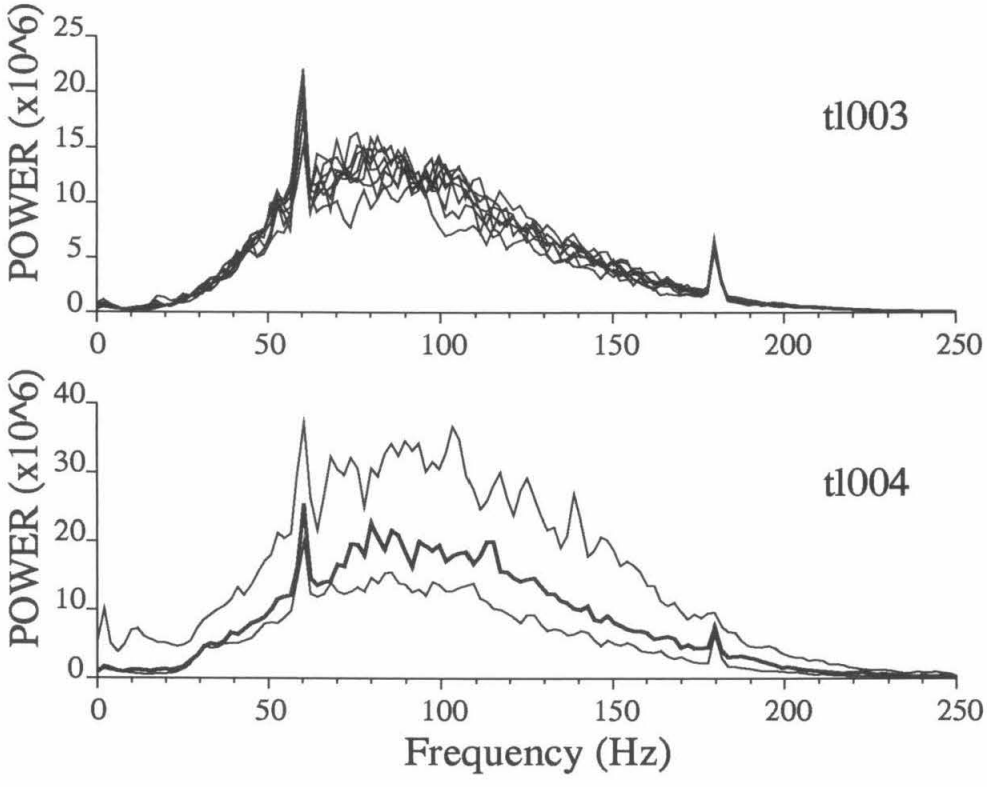

Figure 9.11: Sites 3 and 4. Power spectra are shown for 8 coherence levels for tl003 and for 3 coherence levels for tl004. For tl004, the thick line indicates $c = 0$, the upper line indicates $c = 1$, and the lower line indicates $c = -1$. For tl004, the temporal sequence included 2 sec of visual stimulation. The 60 and 180 Hz interference are prominent, and there is no sign of a peak near 40 Hz.

9.7 Discussion

This preliminary study suggests that there is stimulus-related information in the LFP signal, but the signal quality was too poor to warrant rigorous analysis of the current data. There seems to be a stimulus-locked oscillation near 50 Hz associated with the onset of the stimulus.

To improve the quality of LFP data, a number of modifications should be made to the recording setup. First, an output should be added to the first stage Bak filter in the rack that bypasses the high pass filter. This output could then be used as the input to the new low pass filter. Also, the low pass filter could be modified to include a variable gain amplifier. Second, shielding should be added to reduce the amount of 60 Hz and other low frequency interference. Once these modifications are made, it would be desirable to measure the transfer function of the recording setup by direct methods and also to determine a calibration for the amplitude of the LFP signal.

Chapter 10

Conclusions

The results of this thesis fall into three main categories: (1) the stimulus-locked temporal modulation in both the spike trains and the psychophysics, (2) the gaps and bursts in the unmodulated response to coherent motion and the correlation of gaps between pairs of neurons, and (3) the power spectrum analysis of burst and non-burst cells. Other minor results come from the ROC analysis of bursts and a preliminary examination of the local field potential. These are summarized in the first section, interpretations are given in the second section, and directions for future research are described in the third section.

10.1 Summary

Through a long path which included psychophysical experiments, analysis of previous electrophysiological data, and computational analysis, we have arrived at a firmer conceptualization of the dynamic random dot stimulus and its influence on neurons and behavior. The stimulus is well approximated in the frequency domain as an oriented, three-dimensional cosine, contributed by the signal dots, superposed on a white-noise (flat, only on average) spectrum, contributed by noise dots and the

interaction between signal and noise dots. This can be pictured as a set of louvers, each an infinite plane, that is open at low speeds and begins to close at higher speeds. Each signal dot pair contributes power across the spectrum, while the incidental spatiotemporal interactions of the dots can contribute power to an isolated portion of the spectrum, and may be the cause of the precise temporal modulation that is observed in the spike trains in response to incoherent motion stimulation.

Because we do not know the exact stimulus patterns that generated the modulated responses, our analysis of the temporal structure is incomplete. We have found that neurons will respond with temporal precision between 5–10 msec to the dynamic dot stimulus, and the modulated patterns do not change over time. The autocorrelation of the PSTHs is consistent with these neurons being described by filters that rarely extend beyond 100 msec.

For coherent motion there is virtually no stimulus-locked temporal modulation. A numerical model suggests that the MT response remains unmodulated while the V1 responses are highly modulated because of evenly distributed and appropriately dense spatial sampling, as specified by the Nyquist theorem. Also, because the stimulus contributes power to the preferred as well as null regions in the frequency domain, it is possible that increases and decreases in the number of dots is normalized out.

In the absence of precise stimulus-locked modulation, excessively long gaps appear in the brisk response to coherent, preferred direction motion. A symmetrical bursting phenomenon occurs in the suppressed response to null direction motion. These are, in a crude sense, stimulus-locked also, but they may arise from intrinsic connections in MT rather than directly from the time-varying stimulus. We have argued that the gaps might be due to inhibition and have shown that they are correlated between nearby neurons, but they might also arise from eye movements that cause motion of the scene on the retina. Data from anesthetized animals could help to rule out such eye movements. There is no way to establish solely from the data analyzed here

whether inhibition is actually at play; future experiments will decide.

This thesis began as an investigation into the possibility that oscillations in the 40 Hz range might be related to the coherence of the stimulus or to the monkey's decision. With the propensity that MT cells have for following the dynamics of the stimulus, it is not clear that they would also be able to lock to an underlying 40 Hz rhythm. Of course, if a neuron does not fire on each cycle of the oscillation and if the firing is only loosely locked to the oscillation, then an oscillatory signal may still exist at a population level. This has been demonstrated in a model by Usher et al., (1994).

The final two chapters have elucidated the relationship between burst firing, the refractory period, and peaks in the power spectrum in the 40 Hz range. We show that a simple model using a refractory period, a Poisson process, and with the addition of burst firing is adequate to account for the observed peaks in the power spectrum. It is the length and abruptness of the apparent refractory period and the width of the bursts which determines the shape and location of the peak.

Additionally, we find that counting bursts as single events is better than counting individual spike in terms of an ideal observer who tries to predict the direction of stimulus motion. Finally, we have examined the local field potential at 5 sites and found that there are stimulus locked modulations near 50 Hz during the transient portion of the response, before directional tuning sets in.

10.2 Philosophy and Interpretation

It is not our philosophy that all neurons in one area are performing the same function and can be averaged across and described adequately by a few statistics. We believe each neuron plays a specific role within the cortical microcircuit, and we wish to understand neurons within that context as we try to deduce more about the context

itself. To do this, we need to be familiar with the various types of neurons and their characteristic behaviors. We also need to work on an area of the brain which is well enough understood that computational models exist to guide our search.

One thing that we have learned by examining the temporal structure of the spike trains is how heterogeneous a population of neurons can be. Area MT is considered to be one of the most homogeneous areas, because 90% of cells are directionally selective, but cortical architecture is known to be diverse and more and more is becoming known about this diversity (Agmon and Connors, 1992).

Many successful neuroscientists downplay the usefulness of a highly computational approach. Computational and analytical methods have played an important role in bringing an understanding to phenomena studied here, in particular, in the modeling of the power spectrum and in the analysis of the stimulus in the frequency domain. The model of spatial integration for an MT cell is still being developed, but we think it will prove invaluable for guiding future experimental design. However, we do believe that extra caution may be advisable when using methods that are not intuitive.

Another valuable lesson that has been learned here is the importance of understanding completely the visual stimulus when analyzing fine temporal structure or variability in spike trains, particularly because visual cortical neurons are exquisitely sensitive to fluctuations in the stimulus of which we are probably not aware (Crick and Koch, 1995). Small changes in the stimulus appear in the temporal structure of the cells and may potentially effect the outcome of analysis. The most natural stimuli may not always be the most revealing, and we may be able to do better than traditionally applied stimuli, especially because it has been realized for many years that seemingly complex effects are exerted from beyond the classical receptive field (Allman et al., 1985).

This study separates stimuli into two major classifications: (1) those that induce time modulation in the neuronal response and (2) those that induce an elevated rate

with no time modulation. Some may argue that this difference is merely a matter of time scale, but it has been a useful distinction here.

At the outset, we believed that cortical neurons were highly noisy and unpredictable beyond their mean rate, but this view has been tempered by the observation that the neurons can reproduce firing patterns precisely from trial to trial over the course of hours. With regard to noise, Kelly (1979) has showed that eye movements change perception of spatial frequencies; this also makes each stimulus unique to some extent and can contribute to what might otherwise be considered biological noise. We must accept that even when we present identical stimuli, the neuron under study may receive different signals, but ones that are not necessarily noisy to the extent that they are different.

10.3 Future Research

Examining temporal modulation without having access to the stimulus has whet our appetite to provide a controlled stimulus and link it directly to the time-varying neuronal spike train. For example, one could provide a coherently moving pattern in accordance with a band-limited white-noise velocity signal and then attempt to reconstruct the velocity signal from the spike train, a method first used by de Ruyter van Steveninck and Bialek (1988) to analyze movement-sensitive neurons in the blowfly and more recently to analyze spike trains from electric fish (Gabbiani and Koch, in press). By varying parameters, such as spatial frequency content, in the stimulus and comparing the quality of the reconstruction and by comparing the information that is reconstructed for similar stimuli in V1, we may be able to understand better the transformations within MT and between V1 and MT. One must be careful with this method because it, like many other methods, assumes that the behavior of the neuron is stationary over the analysis window and, in practice, requires the experimenter to

choose the stimulus signal that will be reconstructed.

We believe that cross-correlation analysis will be critical to gain more insight to the function of neurons, and believe that it will be particularly useful in area MT because of the foundations that exist on which to base predictive models. If cells very reliably transform their synaptic inputs to output spike trains, then knowledge of the visual stimulus alone would not allow a complete characterization of the cell's function in conditions where a substantial non-stimulus locked synaptic contribution came from nearby cells, for example, from inhibitory interneurons. As two cells are modulated together, the correlation in the modulation may reveal more subtle differences than that revealed by stimuli that induce no patterning. Many studies concentrate only on the stimulus independent portion of the correlation, i.e., the "shift predictor" is subtracted. We believe that a comparison of the stimulus dependent and the stimulus independent portion of the correlation may prove interesting. In particular, we would like to know which provides more predictive information regarding the spike train of a cortical neuron, the visual stimulus or the spike train from another cortical neuron? We would like to know how the answer varies across pairs of neurons and across stimuli.

With respect to cross-correlation, we also want to develop a quantitative comparison between the effect of common excitatory input which produces narrow peaks, typically less than 10 msec wide, and the effect of correlated gaps, which leads to peaks that are typically 100 msec wide in the interval cross-correlations discussed in Chapter 5.

Our hope for the future is that stimuli can be designed to control the firing of neurons in even more precise ways that will help reveal the functional architecture of the visual cortex. We envision developing a set of stimuli that would allow the characterization of a cell as inhibitory or excitatory, just from the neuronal firing pattern. Work along these lines has already been done *in vitro* (Connors and Gutnick,

1990). The ability of the experimenter to control the firing of the cell with great precision may be an important tool in testing the connections of the cortex. We have tried to demonstrate, based on Fig. 5.6, that imprinting a known pattern in the spike train may be useful.

One part of the temporal analysis that has not been treated here and which we believe holds great promise is the transient part of the response, typically 0–400 msec. This part is highly variable across neurons, sometimes consisting of a slow onset, a sharp peak, or two or three sharp peaks. This temporal structure of the on-transients must arise from the setting up of the cortical network to process the stimulus and may depend on effects from outside the outside the classical receptive field in potentially non-linear ways. The initial observations in local field potentials (Chapter 8.3) are grounds for further investigation. One model has been proposed that attempts to interpret the initial transients in terms of acceleration information (Lisberger and Movshon, 1994).

Many computational models utilize units called "neurons." If these artificial units are individually endowed with computational capabilities similar to their neurobiological counterparts, and if they are connected in networks with similar configurations as, for example, neocortex, then we should expect to see similar behavior on a fine time scale, i.e., at the millisecond time scale. We believe that continued study of fine temporal structure in spike trains will reveal the details of how neurons are connected and how they acquire their tuning properties.

Chapter 11

References

Abeles M (1982) Quantification, smoothing, and confidence limits for single-units' histograms. *J Neurosci Meth* **5:**317–325.

Abeles M (1990) Corticonics. Cambridge: Cambridge UP.

Abeles M, Bergman H, Margalit E, Vaadia E (1993) Spatiotemporal firing patterns in the frontal cortex of behaving monkeys. *J Neurophysiol* **70:**1629–1638.

Adelson EH, Bergen JR (1985) Spatiotemporal energy models for the perception of motion. *J Opt Soc Am A* **2:**284–299.

Adrian E (1928) The basis of sensation: the action of sense organs. London: Christophers.

Aertsen AMHJ, Gerstein GL, Habib MK, Palm G (1989) Dynamics of neuronal firing correlation: Modulation of "Effective Connectivity." *J Neurophysiol* **61:**900–917.

Agmon A, Connors BW (1992) Correlation between intrinsic firing patterns and thalamocortical synaptic responses of neurons in mouse barrel cortex. *J Neurosci* **12:**319–329.

Albright TD, Desimone R (1987) Local precision of visuotopic organization in the middle temporal area (MT) of the macaque. *Exp Brain Res* **65:**582–592.

Allen C, Stevens CF (1994) An evaluation of causes for unreliability of synaptic transmission. *PNAS US* **91:**10380–10383.

Allman J, Miezin F, McGuinness E (1985) Stimulus specific responses from beyond the classical receptive field: neurophysiological mechanisms for local-global comparisons in visual neurons. *Ann Rev Neurosci* **8:**407–430.

Andersen P (1990) Synaptic integration in hippocampal CA1 pyramids. *Prog Brain Res* **83:**215–222.

Bair W, Koch C, Newsome W, Britten K (1993) Temporal structure of spike trains from MT neurons in the awake monkey. Chapter 75, Computation and Neural Systems, Eeckman FH and Bower JM, eds, Kluwer Academic Publishers, Norwell, MA, 495–502.

Bair W, Koch C, Newsome WT, Britten KH (1994) Power spectrum analysis of bursting cells in area MT in the behaving monkey. *J Neurosci* **14:**2870–2892.

Bair W, Koch C (1995) Precision and reliability of neocortical spike trains in the behaving monkey. Chapter 9 in "The Neurobiology of Computation," Bower J, ed. pp. 53–58. Kluwer Academic Publishers, Boston.

Baranyi A, Szente MB, Woody CD (1993) Electrophysiological characterization of different types of neurons recorded in-vivo in the motor cortex of the cat i. patterns of firing activity and synaptic responses. *J Neurophysiol* **69:**1850–1864.

Barlow HB (1972) Single units and sensation: a neuron doctrine for perceptual psychology? *Perception* **1:**371–394.

Barlow HB, Kaushal TP, Hawken M, Parker AJ (1987) Human contrast discrimination and the threshold of cortical-neurons. *J Opt Soc A* **4:**2366–2371.

Bell AJ, Mainen ZF, Tsodyks M, Sejnowski TJ (1995) 'Balancing' of conductances may explain irregular cortical spiking. Technical Report No. INC-9502, Institute for Neural Computation, UCSD, San Diego.

Berman NJ, Douglas RJ, Martin KAC (1992) GABA-mediated inhibition in the neu-

ral networks of visual-cortex. *Prog Brain Res* **90:**443–476.

Beutler FJ, Leneman OAZ (1968) The spectral analysis of impulse processes. *Information and Control* **12:**236–258.

Bialek W, Rieke F, de Ruyter van Steveninck RR, Warland D (1991) Reading a neural code. *Science* **252:**1854–1857.

Bonds AB (1992) Dual inhibitory mechanisms for definition of receptive field characteristics in cat striate cortex. In: *Advances in neural information processing systems*, Vol 4, Moody JE, Hanson SJ, Lippmann RP, eds, pp 75–82. San Mateo, CA: Morgan Kaufmann.

Bracewell RN (1978) The Fourier Transform and its Applications. New York: McGraw-Hill, Inc.

Bradley A, Skottun BC, Ohzawa I, Sclar G, Freeman RD (1987) Visual orientation and spatial frequency discrimination: a comparison of single neurons and behavior. *J Neurophysiol* **57:**755-772.

Britten KH, Shadlen MN, Newsome WT, Movshon JA (1992) The analysis of visual motion: a comparison of neuronal and psychophysical performance. *J Neurosci* **12:**4745–4765.

Britten KH, Shadlen MN, Newsome WT, Movshon JA (1993) Response of neurons in macaque MT to stochastic motion signals. *Visual Neurosci* **10:**1157–1169.

Britten KH, Newsome WT, Shadlen MN, Celebrini S, Movshon JA (1995) A relationship between behavioral choice and the visual responses of neurons in macaque MT. *Visual Neuroscience.* (in press)

Burr DC (1979) Acuity for apparent vernier offset. *Vision Res* **19:**835–837.

Cattaneo A, Maffei L, Morrone C (1981a) Two firing patterns in the discharge of complex cells encoding different attributes of the visual stimulus. *Exp Brain Res* **43:**115–118.

Cattaneo A, Maffei L, Morrone C (1981b) Patterns in the discharge of simple and

complex visual cortical cells. *Proc R Soc Lond B* **212:**279–297.

Champeney DC (1973) Fourier transforms and their physical applications. New York: Academic Press Inc.

Chung SH, Raymond SA, Lettvin JY (1970) Multiple meaning in single visual units. *Brain, Behavior Evol* **3:**72–101.

Connors BW, Gutnick MJ (1990) Intrinsic firing patterns of diverse neocortical neurons. *Trends Neur* **13:**99–104.

Crick F (1984) Function of the thalamus reticular complex: the searchlight hypothesis. *Proc Natl Acad Sci USA* **81:**4586–4590.

Crick F, Koch C (1990) Towards a neurobiological theory of consciousness. *Seminars Neurosci* **2:**263–275.

Crick F, Koch C (1992) The problem of consciousness. *Scientific American* **267:**152–159.

Crick F, Koch C (1995) Are we aware of neural activity in primary visual cortex. *Nature* **375:**121–123.

de Ruyter van Steveninck R, Bialek W (1988) Real-time performance of a movement-sensitive neuron in the blowfly visual system: coding and information transfer in short spike sequences. *Proc R Soc Lond B* **234:**379–414.

Dean AF (1981) The variability of discharge of simple cells in the cat striate cortex. *Exp Brain Res* **44:**437-440.

Derrington AM, Lennie P (1984) Spatial and temporal contrast sensitivities of neurones in lateral geniculate nucleus of macaque. *J Physiol* **357:**219–240.

Eckhorn R, Bauer R, Jordan W, Brosch M, Kruse W, Munk M, Reitboeck HJ (1988) Coherent oscillations: a mechanism of feature linking in the visual cortex? *Biol Cybern* **60:**121–130.

Emerson RC, Bergen JR, Adelson EH (1992) Directionally selective complex cells and the computation of motion energy in cat visual cortex. *Vision Res* **32:**203–218.

Eskandar EN, Richmond BJ and Optican LM (1992) Role of inferior temporal neurons in visual memory: I. Temporal encoding of information about visual images, recalled images and behavioral context. *J Neurophysiol* **68:**1277–1295.

Eysel UT, Burandt U (1984) Fluorescent tube light evokes flicker responses in visual neurons. *Vision Res* **24:**943–948.

Feller W (1968) An introduction to probability theory and its applications, Vol. 1, 3rd Edition. New York: John Wiley & Sons.

Foster KH, Gaska JP, Nagler M, Pollen DA (1985) Spatial and temporal frequency selectivity of neurones in visual cortical areas V1 and V2 of the macaque monkey. *J Physiol* **365:**331–363.

Franklin J, Bair W (1995) The effect of a refractory period on the power spectrum of neuronal discharge. *SIAM J Appl Math* **55:**1074–1093.

Freund TF, Martin KAC, Whitteridge D (1985) innervation of cat visual area-17 and area-18 by physiologically identified X-type and Y-type thalamic afferents. 1. Arborization patterns and quantitative distribution of postsynaptic elements. *J Comp Neur* **242:**263–274.

Gattass R, Gross CG (1981) Visual topography of striate projection zone in posterior superior temporal sulcus (MT) of the macaque. *J Neurophysiol* **46:**621–638.

Gawne TJ, Richmond BJ (1993) How independent are the messages carried by adjacent inferior temporal cortical neurons? *J Neurosci* **13:**2758–2771.

Ghose GM, Freeman RD (1992) Oscillatory discharge in the visual system: does it have a functional role? *J Neurophysiol* **68:**1558–1574.

Gray CM, Singer W (1989) Stimulus-specific neuronal oscillations in orientation columns of cat visual cortex. *Proc Natl Acad Sci USA* **86:**1698–1702.

Gray CM, König P, Engel AK, Singer W (1989) Oscillatory responses in cat visual cortex exhibit inter-columnar synchronization which reflects global stimulus properties. *Nature* **338:**334–337.

Grzywacz NM, Yuille AL (1990) A model for the estimate of local image velocity by cells in the visual cortex. *Proc R Soc Lond B* **239:**129–161.

Gulyas AI, Miles R, Sik A, Toth K, Tamamaki N, Freund TF (1993) Hippocampal pyramidal cells excite inhibitory neurons through a single release site. *Nature* **366:**683–687.

Gulyas B, Orban GA, Duysens J, Maes H (1987) The suppressive influence of moving textured backgrounds on responses of cat striate neurons to moving bars. *J Neurophysiol* **57:**1767–1791.

Hammond P, and MacKay DM (1977) Differential responsiveness of simple and complex cells in cat striate cortex to visual texture. *Exp Brain Res* **30:**275–296.

Han ZS (1994) Electrophysiological and morphological differentiation of chandelier and basket cells in the rat hippocampal formation: a study combining intracellular recording an intracellular staining with biocytin. *Neurosci Res* **19:**101–110.

Harris FJ (1978) On the use of windows for harmonic analysis with the discrete Fourier transform. *Proc IEEE* **66:**51–84.

Heeger DJ (1987) Model for the extraction of image flow. *J Opt Soc Am A* **4:**1455–1471.

Heiden C (1969) Power spectrum of stochastic pulse sequences with correlation between the pulse parameters. *Physical Review* **188:**319–326.

Holt GR, Softky WR, Koch CK, Douglas RJ (1995) A comparison of discharge variability *in vitro* and *in vivo* in cat visual cortex. Submitted for publication.

Kasper E, Larkman A, Blakemore C, Judge S (1991) Physiology and morphology of identified projection neurons in rat visual cortex studied in vitro. *Soc Neurosci Abstr* **17:**114.

Kelly DH (1979) Motion and vision. II. Stabilized spatio-temporal threshold surface. *J Opt Soc Am* **69:**1340–1349.

Kisvarday ZF, Eysel UT (1993) Functional and structural topography of horizontal

inhibitory connections in cat visual cortex. *Eur J Neurosci* **5**:1558–1572.

Knierim J, van Essen D (1992) Neuronal responses to static textural patterns in area V1 of the alert macaque monkey. *J Neurophysiol* **67**:961–980.

Koch C, Crick F (1994) Some further ideas regarding the neuronal basis of awareness. In: *Large-scale neuronal theories of the brain*, Koch C, Davis J, eds, pp 93–109. Cambridge: MIT Press.

Kreiter AK, Singer W (1992) Oscillatory neuronal responses in the visual cortex of the awake macaque monkey. *Eur J Neurosci* **4**:369–375.

Lagae L, Raiguel S, Orban GA (1993) Speed and direction selectivity of macaque middle temporal neurons. *J Neurophysiol* **69**:19–39.

Lee BB, Martin PR, Valberg A (1989) Sensitivity of macaque retinal ganglion cells to chromatic and luminance flicker. *J Physiol* **414**:223–243.

Legéndy CR, Salcman M (1985) Bursts and recurrences of bursts in the spike trains of spontaneously active striate cortex neurons. *J Neurophysiol* **4**:926–939.

Lestienne R, Strehler BL (1987) Time structure and stimulus dependence of precisely replicating patterns present in monkey cortical neuronal spike trains. *Brain Res* **437**:214–238.

Lettvin JP, Maturana HR, McCulloch WS, Pitts WH (1959) What the frog's eye tells the frog's brain. *Proc Inst Rad Eng* **47**:1950–1961.

Levitt JB, Kiper DC, Movshon JA (1994) Receptive fields and functional architecture of macaque V2. *J Neurophys* **71**:2517–2542.

Lisberger SG, Movshon JA (1994) A different approach to modelling pursuit eye movements. In: *Contemporary ocular motor and vestibular research: a tribute to David A. Robinson*, Fuchs AF, Brandt T, Büttner U, Zee D, eds, pp 304–311. Stuttgart: Georg Thieme Verlag.

Lukes T (1961) The statistical properties of sequences of stochastic pulses. *Proc Phys Soc Lond* **78**:153–168.

Lytton WW, Sejnowski TJ (1991) Simulations of cortical pyramidal neurons synchronized by inhibitory interneurons. *J Neurophysiol* **66:**1059–1079.

Magleby KL (1987) Short-term changes in synaptic efficacy. In: *Synaptic function*, Edelman GM, Gall WE, Cowan WM, eds, pp 21–56. New York: John Wiley.

Mainen ZF, Sejnowski TJ (1995) Reliability of spike timing in neocortical neurons. *Science* **268:**1503–1506.

Maunsell JHR, Gibson JR (1992) Visual response latencies in striate cortex of the macaque monkey. *J Neurophysiol* **68:**1332–1344.

Maunsell JHR, Newsome WT (1987) Visual processing in monkey extrastriate cortex. *Annu Rev Neurosci* **10:**363–401.

Maunsell JHR, Van Essen D (1983) Functional properties of neurons in middle temporal visual area of the macaque monkey. I. Selectivity for stimulus direction, speed and orientation. *J Neurophysiol* **49:**1127–1147.

McCormick DA, Connors BW, Lighthall JA, Prince DA (1985) Comparative electrophysiology of pyramidal and sparsely spiny stellate neurons of the neocortex. *J Neurophysiol* **54:**782–806.

McLean J, Palmer LA (1989) Contribution of linear spatiotemporal receptive field structure to velocity selectivity of simple cells in area 17 of cat. *Vision Res* **29:**675–679.

Mikami A, Newsome WT, Wurtz RH (1986) Motion selectivity in macaque visual cortex. I. Mechanisms of direction and speed selectivity in Extrastriate Area MT. *J Neurophysiol* **55:**1308–1327.

Mikami A (1992) Spatiotemporal characteristics of direction-selective neurons in the middle temporal visual area of the macaque monkey. *Exp Br Res* **90:**40–46.

Mitzdorf U (1985) Current source-density method and application in cat cerebral-cortex - investigation of evoked-potentials and eeg phenomena. *Physiol Rev* **65:**37-100.

Mitzdorf U, Li BH, Poppel E (1994) Mass-action view of single-cell responses to stimulation of the receptive-field and/or beyond - exemplification with data from the rabbit primary visual-cortex. *Electroenceph and Clin Neurophys* **92:**442-455.

Morgan MJ, Ward R (1980) Conditions for motion flow in dynamic visual noise. *Vision Res* **20:**431–435.

Newsome WT, Pare EB (1988) A selective impairment of motion perception following lesions of the middle temporal visual area (MT). *J Neurosci* **8:**2201–2211.

Newsome WT, Britten KH, Movshon JA (1989a) Neuronal correlates of a perceptual decision. *Nature* **341:**52–54.

Newsome WT, Britten KH, Movshon JA, Shadlen M (1989b) In: *Proceedings of the retina research foundation, Neural mechanisms of visual perception*, Lam DM-K, Gilbert C, eds, pp 171–198. The Woodlands, TX: Portfolio.

Nowlan SJ, Sejnowski TJ (1994) Filter selection model for motion segmentation and velocity integration. *J Opt Soc Am A* **11:**3177–3200.

Nowlan SJ, Sejnowski TJ (1995) A selection model for motion processing in area MT of primates. *J Neurosci* **15:**1195–1214.

Optican LM and Richmond BJ (1987) Temporal encoding of two-dimensional patterns by single units in primate inferior temporal cortex. III Information theoretic analysis. *J Neurophysiol* **57:**162–178.

Parker A, Hawken M (1985) Capabilities of monkey cortical cells in spatial-resolution tasks. *J Opt Soc Am A* **2:**1101–1114.

Perkel DH, Gerstein GL, Moore GP (1967) Neuronal spike trains and stochastic point processes I, the single spike train. *Biophys J* **7:**391–418.

Poggio GF, Viernstein LJ (1964) Time series analysis of impulse sequences of thalamic somatic sensory neurons. *J Neurophysiol* **27:**517–545.

Press HP, Flannery BP, Teukolsky SA, Vetterling WT (1988) Numerical recipes in C, the art of scientific computing. Cambridge: Cambridge University Press.

Qian N, Andersen RA, Adelson EH (1994) Transparent motion perception as detection of unbalanced motion signals .3. Modeling. *J Neurosc* **14:**7381–7392.

Reichardt, W (1961) Autocorrelation, a principle for the evaluation of sensory information by the central nervous system. In: *Sensory Communication*, Rosenblith WA, ed. New York: Wiley.

Richmond BJ, Optican LM, Podell M, Spitzer H (1987) Temporal encoding of two-dimensional patterns by single units in primate inferior temporal cortex. I. Response characteristics. *J Neurophysiol* **57:**132–146.

Richmond BJ, Optican LM, Spitzer H (1990) Temporal encoding of two-dimensional patterns by single units in primate primary visual cortex. I. Stimulus-response relations. *J Neurophysiol* **64:**351–369.

Richmond BJ, Optican LM (1992) The structure and interpretation of neuronal codes in the visual system. In: *Neural Networks for Perception*, Wechsler H, ed, pp 104–119. New York: Academic Press.

Ridder WH, Tomlinson A (1993) Suppression of contrasts sensitivity during eyelid blinks. *Vision Res* **33:**1795–1802.

Robinson DA (1963) A method of measuring eye movement using a scleral search coil in a magnetic field. *IEEE Trans Biomed Eng* **10:**137–145.

Robson JG (1966) Spatial and temporal contrast-sensitivity functions of the visual system. *J Opt Soc Am* **56:**1141–1142.

Rodman HR, Albright TD (1987) Coding of visual stimulus velocity in area MT of the macaque. *Vision Res* **27:**2035–2048.

Salzman CD, Murasugi CM, Britten KH, Newsome WT (1992) Microstimulation in visual area MT: effects on direction discrimination performance. *J Neurosci* **12:**2331–2355.

Schild D, Schultens HA (1986) The Fourier transform of a peristimulus time histogram can lead to erroneous results. *Brain Research* **369:**353–355.

Sclar G, Maunsell JHR, Lennie P (1990) Coding of image contrast in central visual pathways of the macaque monkey. *Vision Res* **30:**1–10.

Sestokas AK, Lehmkuhle S (1986) Visual response latency of X- and Y-cells in the dorsal lateral geniculate nucleus of the cat. *Vision Res* **26:**1041–1054.

Shadlen MN, Newsome WT (1994) Noise, neural codes and cortical organization. *Curr Opin Neurobiol* **4:**569–579.

Singer W (1994) Putative functions of temporal correlations in neocortical processing. In: *Large-scale neuronal theories of the brain*, Koch C, Davis J, eds, pp 201–237. Cambridge: MIT Press.

Smetters DK, Nelson SB (1993) Estimates of functional synaptic convergence in rat and cat visual cortical neurons. *Soc Neurosci Abstr* **19:**628.

Smith DR, Smith DK (1965) A statistical analysis of the continual activity of single cortical neurones in the cat unanaesthetized isolated forebrain. *Biophys J* **5:**47–74.

Snowden RJ, Treue S, Andersen RA (1992) The response of neurons in areas V1 and MT of the alert rhesus monkey to moving random dot patterns. *Exp Brain Res* **88:**389–400.

Snowden RJ, Treue S, Erickson RG, Andersen RA (1991) The response of area MT and V1 neurons to transparent motion. *J Neurosci* **11:**2768–2785.

Softky WR, Koch C (1992) Cortical cells should fire regularly, but do not. *Neural Computation* **4:**643–646.

Softky WR, Koch C (1993) The highly irregular firing of cortical cells is inconsistent with temporal integration of random EPSPs. *J Neurosci* **13:**334–350.

Somogyi P, Freund TF, Cowey A (1982) The axo-axonic interneuron in the cerebral cortex of the rat, cat and monkey. *Neuroscience* **7:**2577–2607.

Stevens CF (1994) What form should a cortical theory take? In: *Large-scale neuronal theories of the brain*, Koch C, Davis JL, eds, pp 239–255. Cambrigde: MIT Press.

Strehler BL, Lestienne R (1986) Evidence on precise time-coded symbols and memory of patterns in monkey cortical neuronal spike trains. *Proc Natl Acad Sci USA* **83:**9812–9816.

Teich MC (1992) Fractal neuronal firing. In: *Single neuron computation neural nets: foundations to applications*, Mckenna T, Davis J, Zornetzer S, eds. Boston: Academic Press, Inc.

Teich MC, Diament P (1980) Relative refractoriness in visual information processing. *Biol Cybern* **38:**187–191.

Teich MC, Matin L, Cantor B (1978) Refractoriness in the maintained discharge of the cat's retinal ganglion cell. *J Opt Soc Am* **68:**386–401.

Thomson AM, Deuchars J, West DC (1993) Large, deep layer pyramid-pyramid single axon EPSPs in slices rat motor cortex display paired-pulse and frequency-dependent depression, mediated presynaptically and self-facilitation, mediated postsynaptically. *J Neurophysiol* **70:**2354–2369.

Tolhurst D, Movshon J, Dean A (1983) The statistical reliability of signals in single neurons in cat and monkey visual cortex. *Vision Res* **23:**775–785.

Tolhurst DJ, Movshon JA, Thompson ID (1981) The dependence of response amplitude and variance of cat visual cortical neurones on stimulus contrast. *Exp Brain Res* **41:**414–419.

Tomko GJ, Crapper DR (1974) Neuronal variability: non-stationary responses to identical visual stimuli. *Brain Research* **79:**405–418.

Ts'o DY, Gilbert CD (1988) The organization of chromatic and spatial interactions in the primate striate cortex. *J Neurosci* **8:**1712–1727.

Turcott RG, Lowen SB, Li E, Johnson DH, Tsuchitani C, Teich MC (1994) A nonstationary Poisson point process describes the sequence of action potentials over long time scales in lateral-superior-olive auditory neurons. *Biol Cybern* **70:**209–217.

Usher M, Stemmler M, Koch C, Olami Z (1994) Network amplification of local fluc-

tuations causes high spike rate variability, fractal firing patterns and oscillatory local field potentials. *Neural Comp* **6**:795–836.

van Santen JPH, Sperling G (1984) Temporal covariance model of human motion perception. *J Opt Soc Am A* **1**:451–473.

Vaina LM, Sundareswaran V, Harris JG (1995) Learning to ignore—psychophysics and computational modeling of fast learning of direction in noisy motion stimuli. *Cog Brain Res* **2**:155–163.

Victor JD, Purpura K, Katz E, Mao BQ (1994) Population encoding of spatial-frequency, orientation, and color in macaque V1. *J Neurophysiol* **72**:2151–2166.

Vogels R, Orban GA (1990) How well do response changes of striate neurons signal differences in orientation—a study in the discriminating monkey. *J Neurosci* **10**:3543–3558.

Volkmann FC, Riggs LA, Ellicott AG, Moore RK (1982) Measurements of visual suppression during opening, closing and blinking of the eyes. *Vision Res* **22**:991–996.

Vogels R, Spileers W, Orban GA (1989) The response variability of striate cortical neurons in the behaving monkey. *Exp Brain Res* **77**:432–436.

Von der Malsburg C (1981) The correlation theory of brain function. Internal report 81–2, Dept. of Neurobiology, MPI for Biophysical Chemistry Göttingen.

Watson AB, Ahumada AJ Jr (1983) A look at motion in the frequency domain. NASA Technical Memorandum 84352. Also in: *Motion: perception and representation*, Tsotsos JK, ed, pp 1–10. New York: Association for Computing Machinery.

Watson AB, Ahumada AJ Jr (1985) Model of human visual-motion sensing. *J Opt Soc Am A* **2**:322–341.

Werner G, Mountcastle VB (1963) The variability of central neural activity in a sensory system, and its implications for the central reflection of sensory events. *J Neurophysiol* **26**:958–977.

Westheimer G, McKee SP (1977) Perception of temporal order in adjacent visual stimuli. *Vision Res* **17:**887–892.

Zeki SM (1974) Functional organization of a visual area in the posterior bank of the superior temporal sulcus of the rhesus monkey. *J Physiol (Lond)* **236:**549–573.

Zeki SM (1993) A vision of the brain. Oxford: Blackwell Scientific Publications.

Zipser D, Kehoe B, Littlewort G, Fuster J (1993) A spiking network model of short-term active memory. *J Neurosci* **13:**3406–3420.

Zohary E, Hillman P, Hochstein S (1990) Time course of perceptual discrimination and single neuron reliability. *Biol Cybern* **62:**475–486.

Zohary E, Shadlen MN, Newsome WT (1994) Correlated neuronal discharge rate and its implications for psychophysical performance. *Nature* **370:**140–143.