Analysis and Modeling of Spike Train Correlations in the Lateral Geniculate Nucleus

Thesis by

Carlos D. Brody

In Partial Fulfillment of the Requirements
for the Degree of
Doctor of Philosophy

California Institute of Technology
Pasadena, California

1998

(Submitted October 7 1997)

Acknowledgements

First of all, I thank John Hopfield, my advisor. His often astonishing capacity to quickly find and go to the root of a problem, without letting any distracting branchlets get in the way, has been a constant inspiration. It will remain an inspiration and a lesson in effective thinking long after I have left his group. I also thank all the members of the group, from each of whom I have learnt much. Sanjoy Mahajan (our knight of the order of magnitude), Sam Roweis and Erik Winfree carefully read and commented on various parts of this manuscript; Dawei Dong, Maneesh Sahani, and K.P. Unnikrishnan helped with enlightening comments and questions; Laura Rodriguez provided good friendship (and often a feeling of being protected) in addition to her official and helpful role as John's secretary. Though we all work on different questions, the always stimulating level of enthusiastic, diverse communication, and the camaraderie, have made the Hopfield group feel much like a wonderful family. More recently, Amy Forth has provided efficient secretarial support. Many past members of the group also contributed to my intellectual formation as a graduate student: most recently, Tom Annau, Ron Benson, Andreas Herz, and Marcus Mitchell. Other members of the CNS community at Caltech, among them Buster Boahen, Mike Lewicki, Bill Press, and Bruno Olshausen, have also been friends as well as colleagues.

I thank Adam Sillito and members of his group, for their hospitality and for generously sharing their data with me even through sometimes thorny arguments as to its interpretation; this thesis would not have been written without them. I also thank José-Manuel Alonso and R. Clay Reid, who shared retinal ganglion cell data with me for use in the modeling part of the thesis.

The support of my friends has been invaluable; el de mi mamá ha sido feliz y desde siempre.

Finally, I thank JK, for her companionship, and during the last few weeks of thesis writing, her patience and understanding.

Abstract

In this thesis I consider the cross-correlation analysis of spike train data, in two parts. In the first part, I consider the question of the proper interpretation of peaks in covariograms: It is known that peaks in the covariogram of the spike trains of two cells are due to covariations, not time-locked to the stimulus, in the responses of the two cells. Such peaks, even when they have widths on the order of 10s of milliseconds, are often interpreted as evidence of spike timing coordination between the two cells. However, there are other ways to covary which can generate very similar peaks. I describe two of them here: (1) covariations, over different trials, in the overall latency of the response; and (2) covariations, over different trials, in the overall excitability (i.e. average firing rate) of the response. I show how each of these leads to a peak in the covariogram, and how to distinguish such peaks from each other and from peaks due to spike timing coordination. In particular, I describe how to separate the excitability and spike timing components of a covariogram when both types of correlations are present.

The second part of the thesis studies the spike train data obtained in multiple-cell recordings in LGN of cat by Adam Sillito and colleagues (Sillito, Jones, Gerstein and West, 1994, *Nature* 369(6480):479–482). Analysis of this data, including the use of some of the novel insights and methods described in the first part of the thesis, shows that (1) the observed correlations between pairs of cells can be well described in terms of covariations in the latencies and excitabilities of the two cells; (2) the correlations have a time scale of 10s of seconds (20–100 s); (3) the correlations are not specific to the orientation of the drifting sine wave gratings used to drive the geniculate cells under study. A computational model, based on Huguenard and McCormick's (1992) model of thalamic cells, shows that covariations in the resting membrane potential of the two covarying cells can lead to covariogram peaks similar to those found in the experimental data. Together with Sillito et al.'s observation that correlations between pairs of cells only occurred when both members of the pair had the same receptive field types (On/Off/X/Y), and the observation that in some of the records, cells not only fail to positively covary but even *anti*-covary, this suggests that Silltio et al.'s data holds evidence for an intriguing and hitherto undocumented phenomenon: the possibly diffuse, but *cell class-specific*, control of resting membrane potential in LGN neurons.

Contents

List of Figures

Chapter 1 Introduction

Due to the difficulty of recording simultaneously from more than one neuron, most neurophysiological studies are single-cell studies. However, the brain has a massively parallel architecture: all neurons in the brain are simultaneously active to some degree. Single cell studies tell us little about possible interactions between neurons. The situation is similar to observing a multivariate stochastic process and trying to characterize it through observing only its univariate marginal distribution functions. There may be many higher order interactions, critical to the function of the system, which are eluding us. For this reason, there has been consistent interest in multiple-cell recordings. Technical advances have allowed multiple-cell recordings to become increasingly frequent.

One of the principal methods used to analyse multiple-cell spike train data is the use of auto- and cross-covariograms (Perkel et al., 1967; Palm et al., 1988; Aertsen et al., 1989). These characterize the second-order statistics of the responses of two neurons. Two neurons which are independent of each other produce a covariogram whose expected value is zero; if there are significant peaks in the covariogram, then the two neurons recorded from were not independent and therefore there was some interaction between them.

The first part of this thesis is concerned with the question of the proper interpretation of covariogram peaks. Chapter 2 describes how different types of correlations may all lead to similar peaks in a covariogram. This makes the interpretation of a covariogram ambiguous. Understanding how each of the types of correlations can lead to a peak allows devising methods to disambiguate the peaks from each other; methods to do this are described in chapter 2. Chapter 2 will be submitted for publication in *Neural Computation*, verbatim.

The second part of the thesis, chapters 3 and 4, is a study of multiple-cell spike trains recorded in the LGN of cats by Adam Sillito, Helen Jones, George Gerstein, and David West. In their original report on this data (Sillito et al., 1994), the authors focused on the presence of significant peaks in the covariograms of pairs of cells. The new analysis done here also focuses on interactions between pairs of cells, as revealed by the use of spike train covariograms, but uses some of the novel insights gained in chapter 2. The new analysis produces a large set of observations which strongly suggest that the data should be interpreted in a way significantly different to that published in the original report: together with the modeling study described in chapter 4, the analysis suggests that the peaks in the covariograms are caused by slow (timescale tens of seconds) covariations in the resting membrane potentials of the two neurons being recorded from.

It is known that various regions in the brain stem (in particular, the parabrachial region) and geniculate-projecting areas of cortex (areas 17, 18, and 19 in the cat; mostly area 17 in the monkey) may influence the resting membrane potential of cells in the LGN through activation of metabotropic receptors. The timescale of their effects is in the tens of seconds, similar to the timescale of interactions found in the analysis done here. Limited evidence available so far has suggested that the control of resting potential is global over the entire LGN (Mukherjee et al., 1995). However, Sillito et al. reported that only pairs of neurons with similar receptive field properties (On/Off/X/Y) were correlated. In all records where three cells were successfully simultaneously recorded from, one of the three cell pairs always covaried much more strongly than the other two cell pairs. Furthermore, while doing the analysis reported on in chapter 3, I observed that for some records, some cell pairs even *anti*-covaried. Thus, the analysis and modeling done in this thesis indicate that Sillito et al.'s data holds evidence for an intriguing and hitherto undocumented phenomenon: the possibly diffuse, but *cell class-specific*, control of resting membrane potential in different groups of LGN neurons.

Chapter 2 Latency, Excitability, and Spike Timing Correlations

Abstract

It is known that peaks in the shuffle-corrected cross-correlogram of the spike trains of two cells are due to covariations, not time-locked to the stimulus, in the responses of the two cells. Such peaks are often interpreted as evidence of spike timing coordination between the two cells. However, there are other ways to covary which can generate very similar peaks, particularly when the peaks have widths on the order of tens of milliseconds. I describe two of these here: (1) covariations, over different trials, in the overall latency of the response; and (2) covariations, over different trials, in the overall excitability (i.e. average firing rate) of the response. I show how each of these leads to a peak in the covariogram, and how to distinguish such peaks from each other and from peaks due to spike timing coordination. In particular, I show how to separate the excitability and spike timing components of a covariogram when both types of correlations are present. In general, it must be noted that shuffle-corrected correlograms – equivalent to covariograms – measure covariations of *any* kind between the firing patterns of two cells. Latency, excitability, and spike timing are simply three examples of particular relevance to neurophysiology.

2.1 Introduction

Imagine that we have obtained spike trains from two neurons, recorded simultaneously during many identically prepared trials of an experiment. To assess whether there are interactions between the two neurons, we build the cross-correlogram of the two spike trains, averaged over trials. We must take into account the fact that the correlogram of two cells with similar peri-stimulus time histograms (PSTHs) will be peaked even if the two cells are independent and there are no interactions between them. This case is analogous to that of two scalar variables, s_1 and s_2, that are independent but have similar means: their correlation $r = \langle s_1 s_2 \rangle$ will be positive even though they are independent. But their covariance $v = \langle (s_1 - \bar{s_1})(s_2 - \bar{s_2}) \rangle$ (where $\bar{s_i} = \langle s_i \rangle$ and $\langle\rangle$ represents averaging over samples) is different from zero only if s_1 and s_2 are not independent, regardless of what their mean values are. The scalar case suggests what to do in the spike train case in order to obtain a straightforward indicator of departure from independence: use a measure analogous to covariance and compute cross-*covariograms* instead of cross-correlograms. (Cross-covariograms, also known in the neurophysiology literature as *shuffle-corrected cross-correlograms* (Perkel et al., 1967; Palm et al., 1988; Aertsen et al., 1989), are described in section 2.2 below. I will sometimes shorten the term "cross-covariogram" to simply "covariogram".) Cross-covariograms are different from zero only if the two spike trains are not independent, regardless of whether their PSTHs are similar or not.[1] How should one interpret the presence of a peak in the covariogram? Figure 2.1 shows how three very different types of spiking patterns can all lead to similar peaks in the covariogram. Distinguishing which of these types of spiking patterns (spike timing, latency, or excitability) caused the peak is important, since that distinction will strongly affect any interpretation of the covariogram. In this paper I will explain how latency and excitability correlations give rise to a peak (spike timing correlations have been treated before, e.g. Perkel et al. 1967), and how to disambiguate the three types of patterns.

In each trial of Figure 2.1A most spikes from cell 1 have a corresponding, similarly timed spike in cell 2. Let us call this type of correlation, where there is coordination between individual spike times in each trial, a *spike timing correlation.* In each trial of Figure 2.1B the two cells spike independently of each other, so spikes in cell 1 do not have any particular corresponding spikes in cell 2. But on each trial the spike trains of both cells have been shifted together in time by some amount that is different from trial to trial. Let us call this type of correlation, where

---

[1] It is sometimes said that the covariogram is a correlogram that has been corrected for "stimulus-locked correlations".

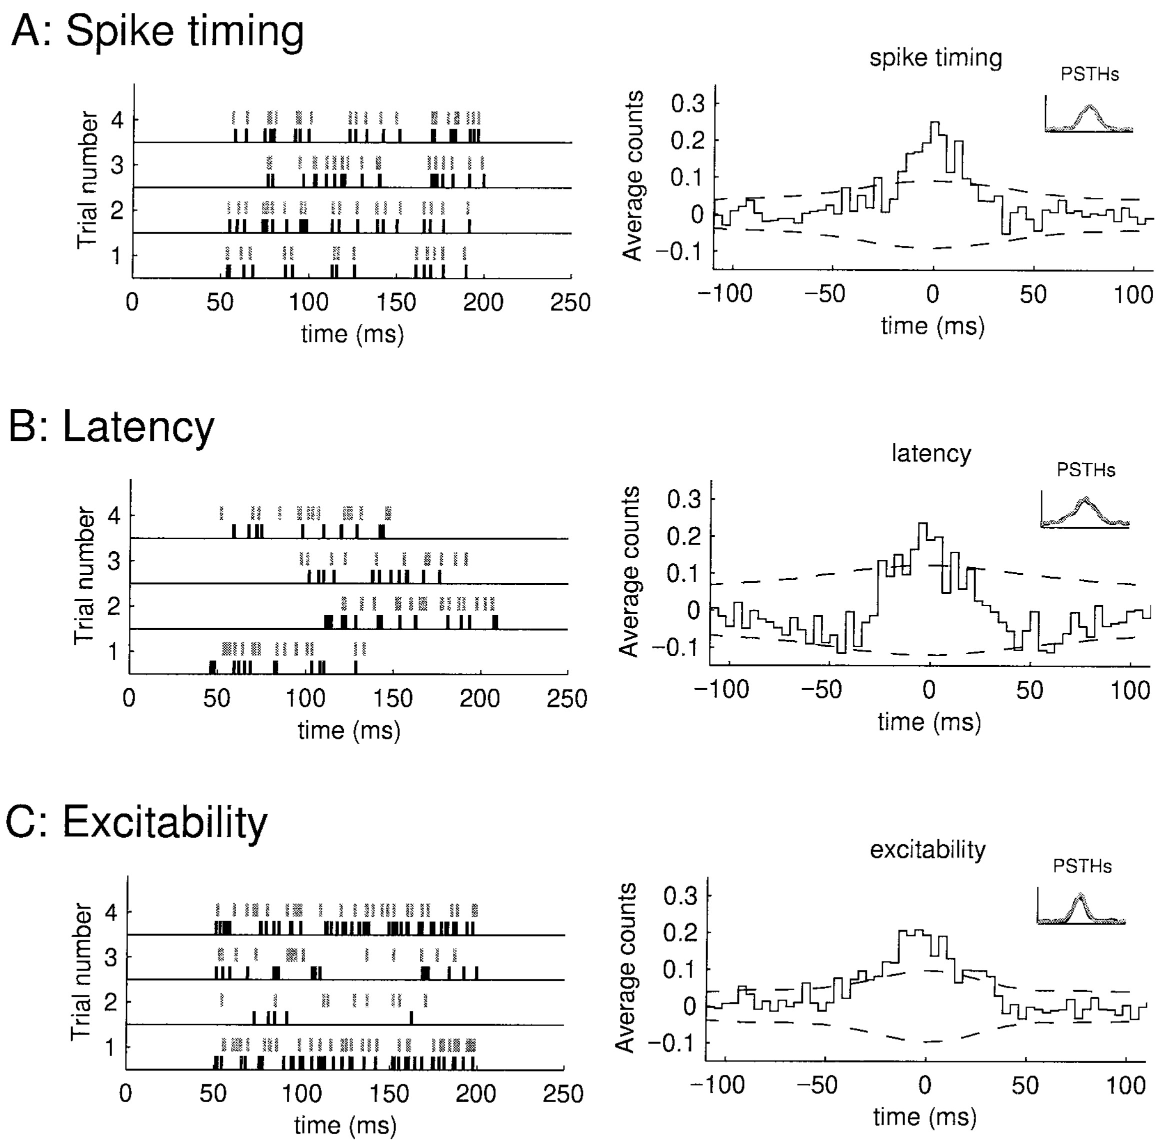

Figure 2.1: Three types of correlations. Despite the very different spiking patterns, all three shuffle-corrected correlograms (henceforth called covariograms– see section 2.2) look similar. Each row illustrates a type of spiking pattern: On the left is a raster plot of two simulated cells. On the right is the covariogram of spike trains made with the same type of spiking pattern illustrated to its left. (Parameters used to generate the spike trains on the left were set to extreme values, to emphasize clarity in illustrating each pattern; parameters used to generate spike trains for the covariograms on the right were set to "physiologically plausible" values.) **A:** On each trial, most spikes in cell 1 have a corresponding, closely timed, spike in cell 2. Both cells have the same response latency and overall firing rate in all trials. **B:** Spikes in cell 1 do not have a corresponding spike in cell 2: on each trial, the two spike trains were generated independently of each other. But the latency of the response varies together over trials. **C:** Again, on each trial the spikes for the two cells were generated independently of each other. But the average firing rate of the response varies together over trials. Zero counts on the covariogram y axes is the expected value if the two cells are independent; the dashed lines are two standard deviations away from this expected value. The inset at the top right of each covariogram shows the PSTHs of the two cells involved, plotted on axes that are 250 ms wide and 60 Hz tall.

the coordination is in the time shifts of the responses of the two cells, a *latency correlation*.[2] In Figure 2.1C, the total number of spikes fired in each trial by cell 1 covaries with the total number of spikes fired in each trial by cell 2 (i.e. the two cells have similar gains in each trial, where "gain" measures how many spikes each neuron produces in response to the stimulus driving it). But within each trial the spike trains of the two cells were generated independently of each other. Let us call this last type of correlation, where the coordination is between the gains of both cells, an *excitability correlation*. In all, Figure 2.1 shows that the covariogram and the PSTHs alone do not provide enough information to conclude which type, or mix of types, of spiking patterns gave rise to the peak in the covariogram.

The three different spiking patterns illustrated in Figure 2.1 can give rise to strikingly similar covariograms because the PSTHs and joint peristimulus time histogram (JPSTH; the cross-correlogram is a summary of the JPSTH) estimate only the first and second-order moments of the distribution of the responses of the cells. These moments only *constrain* what the distribution can be, but do not specify it uniquely. Some possibilities may be eliminated as incompatible with known physiology; others, particularly those illustrated in this paper, are relevant to neurophysiological studies and need to be considered further. In the paper, I will make no assumptions about the mechanisms that might have generated covariations in latency or excitability, so that the concepts to be explained will apply equally well to any two simultaneously recorded point processes: different types of spiking patterns will be illustrated using rasters made by simulating Poisson processes with time-varying event densities (Papoulis 1984). But the language used to illustrate the ideas will be that of neurophysiology ("cells", "spikes", etc.), and the neurophysiological relevance of the concepts and possible physiological mechanisms involved in generating the patterns will also be discussed.

Section **2.2** introduces the notation and correlogram methods used in the paper. Section **2.3** describes how each of the spiking patterns illustrated in Figure 2.1 leads to a peak in the covariogram. Section **2.4** presents methods to disambiguate the three patterns. Finally, section **2.5** discusses the results and their relevance to neurophysiology. A preliminary version of these results has appeared in abstract form (Brody, 1997).

---

[2] The word "latency" here indicates the time shift of the whole response, not just of the first spike.

2.2 Notation and correlogram methods

The spike trains of two cells will be represented by two time-dependent functions, $S_1(t)$ and $S_2(t)$. They will be assumed *binned* [3] (thus t will index a discrete time) and collected over many experimental trials, indexed by a superscript r. The cross-correlogram of each trial is defined as

$$C^r(t) \;\equiv\; \sum_\tau S_1^r(\tau + t) S_2^r(\tau) \;\equiv\; S_1^r \otimes S_2^r. \tag{2.1}$$

In a manner exactly analogous to estimating the covariance of two scalar stochastic variables s_1 and s_2 by $\langle\, (s_1 - \bar{s}_1) \cdot (s_2 - \bar{s}_2)\, \rangle$, where $\bar{s}_i = \langle s_i \rangle$ and $\langle\rangle$ stands for an average over trials r, the estimator of the cross-covariogram of S_1^r and S_2^r is

$$\begin{aligned} V &= \langle\, (S_1^r - \overline{S_1^r}) \otimes (S_2^r - \overline{S_2^r})\, \rangle \\ &= \langle S_1^r \otimes S_2^r \rangle - \langle \overline{S_1^r} \otimes S_2^r \rangle - \langle S_1^r \otimes \overline{S_2^r} \rangle + \overline{S_1^r} \otimes \overline{S_2^r} \\ &= \langle S_1^r \otimes S_2^r \rangle - \langle S_1^r \rangle \otimes \langle S_2^r \rangle. \end{aligned} \tag{2.2}$$

When trials are time-locked to a stimulus, $\langle S_1^r \rangle$ is called the *peri-stimulus time histogram* (PSTH) of cell 1, and $\langle S_2^r \rangle$ is the PSTH of cell 2. The two terms in equation (2.2) are the *raw cross-correlogram* $R = \langle S_1^r \otimes S_2^r \rangle$ and the *shuffle corrector* $K = \langle S_1^r \rangle \otimes \langle S_2^r \rangle$. The cross-covariogram $V = R - K$ is also known as the *shuffle-corrected correlogram* (Perkel et al., 1967; Palm et al., 1988; Aertsen et al., 1989); however, I will call it here the *cross-covariogram*, or sometimes simply the *covariogram*, for short.[4]

Regardless of whether the two PSTHs are similar or not, if the two spike trains S_1^r and S_2^r are

---

[3] We will be interested in the time density of spiking events; if individual spiking events are represented by Dirac delta functions, binning serves the role of smoothing the train of delta functions which represents each spike train. I have chosen binning (equivalent to smoothing with a boxcar function and sampling at a rate inversely proportional to the width of the boxcar) over smoothing with a Gaussian function because the calculation of significance limits (equation (2.3) below) is computationally simpler in the binning case. To obtain some of their results, Palm et al. (1988) assumed that the bin width was small enough that no single trial had more than one spike in any bin; this assumption will not be used here because there are cases where it would require bins too small to be effective in smoothing. For example, bursty neurons can fire at up to 800 Hz (Gray and Mccormick, 1996), which would require bin widths of less than 1.25 ms. In contrast, typical bin widths used here are in the range of 4 to 8 ms. For this reason, significance will not be quantified using the "surprise" measure of Palm et al. (1988), but will instead be quantified using equation (2.3).

[4] The *shift predictor* (Perkel et al., 1967) is very similar to the shuffle corrector, except that K is replaced by $P = \langle S_1^r \otimes S_2^{\Pi(r)} \rangle$, where $\Pi(r)$ is some permutation of the stimulus presentations r (and the corresponding substitution is made in equation (2.2)). If different trials are independent of one another, then the expected value of the shift predictor P is equal to the expected value of the shuffle corrector K. Thus they are both estimators of the same function. In practice it is preferable to use K instead of P since the former is a less noisy estimator: K can be written as the average of P, taken over the set of all possible permutations Π (Palm et al., 1988).

independent then the expected value of V, where the expectation is taken over repetitions of an entire multi-trial experiment, is zero:

$$\begin{aligned} E\{V\} = E\{R-K\} &= E\{\langle S_1^r \otimes S_2^r \rangle\} - E\{\langle S_1^r \rangle \otimes \langle S_2^r \rangle\} \\ &= E\{S_1^r\} \otimes E\{S_2^r\} - E\{S_1^r\} \otimes E\{S_2^r\} \;=\; 0. \end{aligned}$$

Significant departures of V from zero indicate that the two cells are not independent. This conclusion does not require any assumptions about the distributions from which S_1^r and S_2^r are drawn. Estimating the *significance* of departures of V from 0 does require some assumptions. For the null hypothesis I will assume that S_1 is independent of S_2, that different trials of S_1 are independent of each other, and that different bins within each trial of S_1 are independent of each other (and I will make similar assumptions for the trials and bins of S_2). Then the variance in the null hypothesis V is

$$\sigma^2(t) = (\sigma_1^2 \otimes \sigma_2^2 + \mu_1^2 \otimes \sigma_2^2 + \sigma_1^2 \otimes \mu_2^2)/N_{\text{trials}}, \tag{2.3}$$

where $\mu_1(t)$ and $\sigma_1^2(t)$ are the mean and variance of $S_1^r(t)$ over trials r, respectively, similarly for μ_2 and σ_2^2, and N_{trials} is the number of trials in the experiment. In practice, one uses the sample means and variances to calculate $\sigma(t)$. The 2σ limits, calculated in this way, are displayed as dashed lines in the covariograms throughout this paper.

The *Joint Peristimulus Time Histogram*, or JPSTH (Aertsen et al., 1989), is a matrix, built out of the cross products of S_1^r and S_2^r. The raw JPSTH, analogous to the raw cross-correlogram R, is $J_R(\tau_1, \tau_2) = \langle S_1^r(\tau_1) S_1^r(\tau_2) \rangle$. The "PST predictor", analogous to the shuffle corrector, is $J_K = \langle S_1^r(\tau_1) \rangle \langle S_2^r(\tau_2) \rangle$. The difference of these two is the corrected JPSTH,

$$J_V(\tau_1, \tau_2) = J_R - J_K = \langle S_1^r(\tau_1) S_2^r(\tau_2) \rangle - \langle S_1^r(\tau_1) \rangle \langle S_2^r(\tau_2) \rangle. \tag{2.4}$$

If S_1 and S_2 are independent, the expected value of J_V is zero. The covariogram V can be obtained from J_V by summing along diagonals of the matrix such that $\tau_1 - \tau_2$ is kept constant. Following Aertsen et al. (1989), a normalised version of J_V is obtained by computing, bin by bin, the correlation coefficients of $S_1^r(\tau_1)$ and $S_2^r(\tau_2)$. In this paper figures involving a JPSTH display the matrix of correlation coefficients

$$J_{\mathrm{norm}}(\tau_1, \tau_2) = J_V(\tau_1, \tau_2) \; / \; (\sigma_1(\tau_1)\sigma_2(\tau_2)) \,, \tag{2.5}$$

where $\sigma_1^2(\tau_1)$ is the variance, over different trials, in $S_1^r(\tau_1)$, and $\sigma_2^2(\tau_2)$ is the variance in $S_2^r(\tau_2)$. The correlation coefficients $J_{\mathrm{norm}}(\tau_1, \tau_2)$ are bounded within $[-1, 1]$. If the coefficient $J_{\mathrm{norm}}(\tau_1, \tau_2) = 1$, then $S_1(\tau_1)$ and $S_2(\tau_2)$ are perfectly correlated: $S_1(\tau_1) = \alpha\, S_2(\tau_2)$ for some positive constant α.

2.3 Different types of correlations

This section examines each of the three types of spiking patterns introduced in section **2.1**. Section **2.4** will describe methods to disambiguate the three patterns.

2.3.1 Excitability correlations

How do excitability correlations lead to a peak in the covariogram? Let n_1^r and n_2^r be the total number of spikes fired in trial r by cell 1 and cell 2, respectively. Then (as shown below) the DC value of the covariogram, $\sum_t V(t)$, is equal to $\langle n_1^r n_2^r \rangle - \langle n_1^r \rangle \langle n_2^r \rangle$, which is the covariance of n_1^r and n_2^r. If both cells have similar gains which vary over trials, then n_1^r and n_2^r will have a positive covariance, so the DC value of the covariogram will be positive. This necessarily implies that the covariogram cannot be zero for all t. Since we are using trials that each last a finite amount of time, the covariogram falls to zero at large $|t|$. Therefore, if the covariogram's DC value is positive, there must be a positive peak somewhere in the covariogram. This is true independently of whether or not there was coordination between individual spike times in the two cells. If there was no coordination between individual spikes times, we will call the peak an excitability correlation peak.

Let us now prove the main statement in the argument above: $\sum_t V(t) = \langle n_1^r n_2^r \rangle - \langle n_1^r \rangle \langle n_2^r \rangle$.

The area under a spike train correlogram is independent of the bin width used to compute it. Letting

$$n_1^r \equiv \sum_t S_1^r(t) \qquad \text{and} \qquad n_2^r \equiv \sum_t S_2^r(t) \tag{2.6}$$

be the number of spikes fired in trial r by cell 1 and cell 2, respectively, the area under the rth trial's raw cross-correlogram, C^r, is

$$\begin{aligned} \sum_t C^r(t) &= \sum_{t=-\infty}^{\infty} \sum_{\tau=-\infty}^{\infty} S_1^r(\tau + t) S_2^r(\tau) = \sum_u \sum_\tau S_1^r(u) S_2^r(\tau) = \sum_u S_1^r(u) \sum_\tau S_2^r(\tau) \\ &= n_1^r n_2^r. \end{aligned} \tag{2.7}$$

Similarly, the area under the shuffle corrector K is always $\langle n_1^r \rangle \langle n_2^r \rangle$, independently of bin width. If we sum across time in the covariogram, then

$$\begin{aligned} \sum_t V(t) &= \sum_t R(t) - K(t) = \sum_t \langle C^r(t) \rangle - \sum_t K(t) = \langle \sum_t C^r(t) \rangle - \sum_t K(t) \\ &= \langle n_1^r n_2^r \rangle - \langle n_1^r \rangle \langle n_2^r \rangle \end{aligned} \tag{2.8}$$

which is the covariance of n_1 and n_2, as claimed above.

If we assume that all trials are of the same length of time, then equation (2.8) is directly proportional to the covariation in the mean firing rates of the two cells. What kind of shape does a peak due to positive excitability correlations have? The width of the peak in the covariogram is directly proportional to the temporal width of the parts of the cells' responses that had comodulated gains. Figure 2.2 shows an example. Here the part with comodulated gains was a roughly bell-shaped firing rate function which peaked at 100 ms and was approximately 60 ms wide. Hence the width of the covariogram peak is on the order of $2 \times 60 = 120$ ms. The DC value of the corrected, but unnormalized, JPSTH is also equal to the covariation in the total number of spikes fired per trial (equation (2.8)). Since normalization changes the magnitude, but not the sign, of each of the elements of the JPSTH matrix the excitability peak can also be seen in the normalized JPSTH.

Covariations in the *number* of spikes fired may, of course, have been accompanied by fine control in the *timing* of the spikes. The DC value of the covariogram does not in itself tell us whether this was the case.[5] To address that question we have to use other methods, explored further in section 2.4.1.

---

[5] For example, the reader may notice that the covariograms in figures 2.1A and 2.4, which were labelled as "spike timing covariograms," have clearly positive DC values. This follows from the way rasters for these covariograms were generated (see appendix 2.7.3). An example of a spike timing correlation *without* an accompanying excitability correlation can be found in the left panel of figure 2.7 (see appendix 2.7.3).

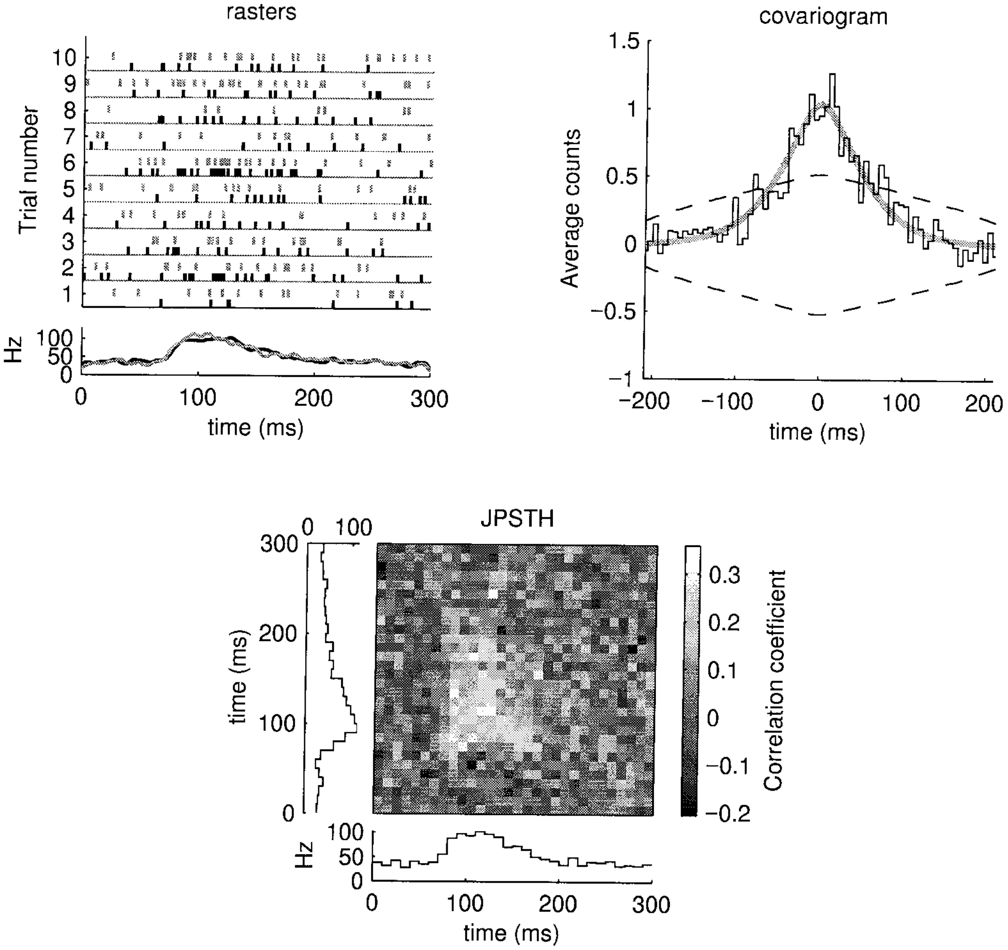

Figure 2.2: Excitability correlations, by construction. The two cells were simulated by Poisson processes with time-varying event densities, to represent time-varying firing rates; event trains for the two cells in each trial were drawn independently from each other. Each cell had two components in its firing rate: a constant background rate of 35 Hz, plus an alpha function shaped (roughly bell-shaped) firing rate function of width ≈ 60 ms. On each trial the height of the alpha function was the same for both cells, but it varied from trial to trial. (Averaged over trials, the peak of the alpha function was at 100 ms and was 70 Hz high.) Since the component of the two firing rates that had covarying gains was approximately 60 ms wide, the peak in the covariogram shown in the top right is on the order of $2 \times 60 = 120$ ms wide. 200 rasters were generated for this covariogram. The thick grey line is the analytical expected value from equation (2.19), computed with full knowledge of the parameters and procedures used to run the stochastic process that generated the spike trains. The two thin dashed lines are significance limits. On the top left are ten of the actual rasters used to compute the covariogram, and below them the PSTH of the two cells (smoothed by a Gaussian with 6 ms standard deviation). On the bottom is the normalised JPSTH, flanked by PSTHs of each of the two cells. See Appendix 2.7.1 for details of how the rasters were constructed and how the analytical curve going through the covariogram was calculated.

2.3.2 Latency correlations

How do latency correlations lead to a peak in the covariogram? Suppose we are recording from two neurons whose responses are independent. In that case the expected cross-covariogram (see equation (2.3)) is zero,

$$\underbrace{E\{V\}}_{\text{covariogram}} = \underbrace{E\{R\}}_{\text{raw x-corr}} - \underbrace{E\{K\}}_{\text{corrector}} = 0,$$

so for a single repetition of the multi-trial experiment, we expect that within sampling noise $K \approx R$. Now imagine that we are recording from two cells whose responses are independent – except for the *latencies* of the responses, to be understood as follows: take the responses obtained from two completely independent cells, and in each trial shift the responses of *both* cells by some amount of time t^r. Let t^r be different for each trial. The resulting spike trains will be responses that we will describe as being independent except for the latencies (i.e. the time shifts) of the responses. How will the time shifts affect V? Let us examine its two components, R and K. The raw cross-correlogram R will not be affected: cross-correlograms depend only on relative times between the two cells, and the relative times between the spikes of the two cells on each trial have not changed, since on each trial the responses of both cells were shifted together. The shuffle corrector K, on the other hand, *will* be affected: $K = \langle S_1 \rangle \otimes \langle S_2 \rangle$, and the two PSTHs $\langle S_1 \rangle$ and $\langle S_2 \rangle$ will each be smoothed by the time jittering caused by the shifts t^r. The shuffle corrector K will be a smoothed version of the shuffle corrector that would have been obtained if there had been no time shifts. The shuffle corrector obtained without time shifts is $K \approx R$; the smoothed version of this that is obtained with time shifts is a smoothed version of R. Thus, if there are latency time shifts, sharp peaks or troughs in R will be smoothed away in the shuffle corrector K, and will remain outstanding in the cross-covariogram $V = R - K$. The top left panel of Figure 2.3 is a schematic illustration of this, showing a raw cross-correlogram R, a smoothed version of the raw cross-correlogram, which represents K, and the difference of the two. This last is the expected covariogram V. (Since the preceding description excluded excitability correlations, the DC value of V is zero. The peak is surrounded by two compensating troughs.) A more mathematically detailed description of latency correlations can be found in Appendix 2.8.

The three other panels of Figure 2.3 show the covariogram, JPSTH, and 10 rasters of spike trains constructed with latency correlations, made by simulating two independent Poisson cells and then shifting both spike trains in each trial by a random amount of time t^r drawn from a

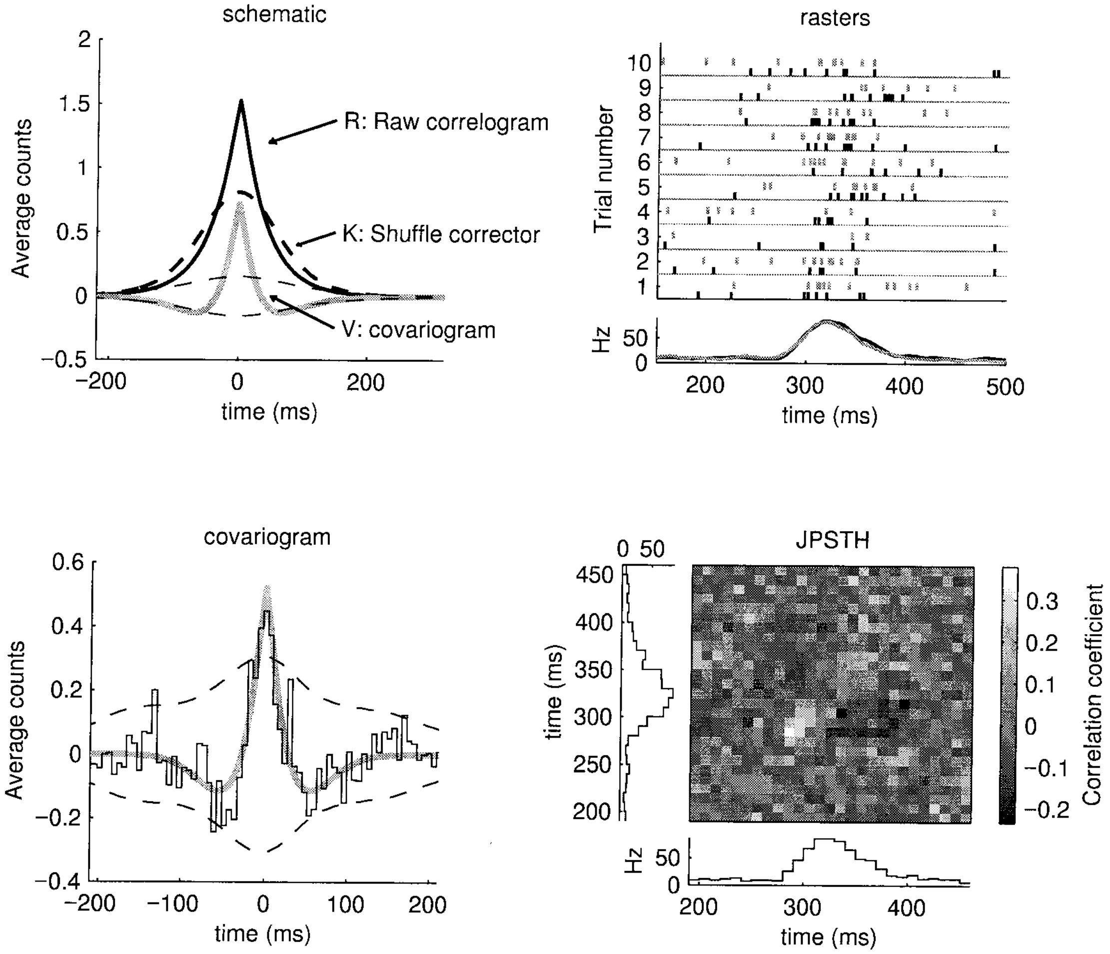

Figure 2.3: Latency correlations. The top left panel shows a schematic of how latency correlations lead to a peaked covariogram (see text for explanation). The three other panels show an example of latency correlations, using rasters made by simulating two independent Poisson cells, each raster pair of which was then shifted in time by a random amount drawn anew for each trial from a Gaussian distribution with mean 0 ms and standard deviation 15 ms. See Appendix 2.7.2 for details. 200 rasters were generated; 10 of the rasters are shown in the top right. Below them are the PSTHs of all 200 rasters for both cells. Overlayed on the covariogram at the bottom left is the analytical expected value of the covariogram from equation (2.20), computed with full knowledge of the parameters and procedures used to run the stochastic process that generated these spike trains. The two thin dashed lines are significance limits. Notice the weak but clearly present troughs in the JPSTH at the bottom right.

Guassian with zero mean and standard deviation 15 ms. The covariogram at the bottom left has the peak and compensating troughs illustrated in the schematic at the top left.

2.3.3 Spike timing correlations

For comparison with latency and excitability correlations, Figure 2.4 shows some example rasters, the covariogram, and the normalised JPSTH of spike trains constructed with spike timing correlations.

These were defined in the introduction as correlations where, in each trial, there is some coordination between *individual* spike times of two cells. A specific example of a mechanism leading to such correlations is that of a fast excitatory synaptic connection from one cell to the other, where there is a (probabilistic) causal relationship between individual presynaptic spikes and one or a few closely timed postsynaptic spikes (or zero spikes if the postsynaptic cell failed to fire). Given arbitrary network topologies, further complications in the connectivity of the cells and the resulting shape of the covariogram are easily imagined. Thus spike timing correlations are very *flexible*. Almost any shape in a covariogram can be explained in terms of spike timing correlations.[6] In contrast, the covariogram shapes that latency or excitability correlations can give rise to are closely tied to their underlying PSTHs and are far less arbitrary. Spike timing correlations are also very *complex*, in the sense that during each trial the two cells share information about a large number of parameters (many individual spike times). This is obvious in the case of a direct connection from one cell to another, but it is also true if the information sharing mechanism is indirect (e.g. there is a common input to both cells). The number of parameters about which information is being shared on each trial is the most important distinction between spike timing correlations and latency or excitability correlations.

2.4 Disambiguating correlations

In order to correctly interpret a covariogram, it is important to distinguish the covariogram shapes due to each type of spiking pattern from each other. This section describes quantitative methods to do so.

---

[6] In particular, spike timing correlations do not necessarily imply a thin covariogram peak. The width of the peak (or trough) in the covariogram will depend on the particular mechanism involved in generating it. Fast ionotropic synaptic mechanisms would, by themselves, lead to thin peaks. But other, slower-acting mechanisms (e.g. NMDA or GABA$_B$ synapses (Brody, 1996), metabotropic receptors, or interactions with slow conductances intrinsic to the cells in question (Kirkland and Gerstein, 1996) would lead to broad peaks.

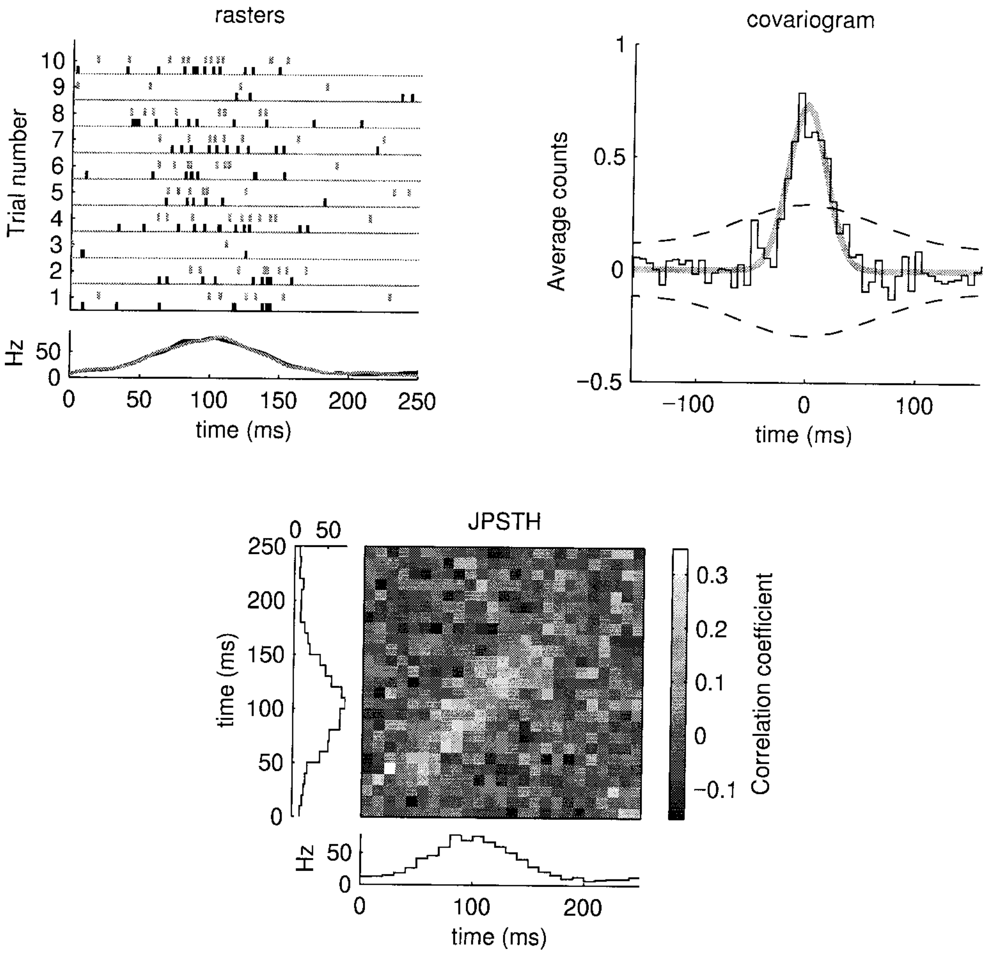

Figure 2.4: Spike timing correlations, by construction: most individual spikes of one cell have a specific, corresponding, closely timed spike in the other cell. On each trial, the corresponding spikes in the two cells were generated by randomly jittering the times of a single spike train twice; the result of the first random jittering was assigned to cell 1, and the result of the second random jittering was assigned cell 2. Jittering was done by adding a random amount, drawn i.i.d. from a Gaussian distribution with mean zero and standard deviation 12 ms, to the time of each spike. The unjittered spike train from which the jittered versions were made was produced anew for each trial. In addition, random background spikes were also added to the trains of the two cells. See Appendix 2.7.3 for details. Ten of the rasters made in this way are shown on the top left. The covariogram of two hundred such rasters is shown on the top right. The thick grey line in the covariogram is the analytical expected value from equation (2.21), calculated with full knowledge of the parameters and procedures used to generate the spike btrains. The normalised JPSTH is shown at the bottom.

2.4.1 Identifying the presence of excitability correlations

Consider two cells producing spikes trains S_1^r and S_2^r. Let n_1^r and n_2^r be the total number of spikes fired in each trial r by cell 1 and cell 2, respectively. Suppose the covariation

$$\nu = \langle n_1^r n_2^r \rangle - \langle n_1^r \rangle \langle n_2^r \rangle \tag{2.9}$$

is positive. This alone implies there is a peak in the covariogram of S_1 and S_2 (see section 2.3.1). But what parts of a possibly multi-peaked covariogram shape will be due to covariations in parameters such as overall gains[7] (that is, excitability correlations), and what parts will be due to other types of interactions (e. g. spike timing or latency correlations)? Based on the experimentally observed variations in the gains of the cells' responses, an estimate of the covariogram that would due to excitability correlations alone can be obtained; subtracting this estimate from the original covariogram will then leave an estimate of the contributions due to other types of correlations.

In making the excitability estimate, it is important to distinguish between changes in the gain of the cells' stimulus-induced responses and changes in the cells' background firing rates. Considering one cell initially, let us first divide its PSTH into two terms:[8] a time-independent background firing rate b, and a time-dependent stimulus-induced firing rate $p(t)$:

$$\mathrm{PSTH}(t) \equiv \langle S^r(t) \rangle = b + p(t). \tag{2.10}$$

Set $b = \langle n_b^r \rangle / N_B$, where n_b^r is the number of experimentally observed spikes that occurred before the stimulus start time (written as t_0) in trial r, and N_B is the number of bins in the PSTH before t_0. The magnitude and time dependence of $p(t)$ are then set by $p(t) = \mathrm{PSTH}(t) - b$. Now let us add a description of the experimentally observed trial-by-trial variations in the gains of b and $p(t)$. To do this, let us model the experimentally observed responses with a (possibly stochastic) model cell which has responses $M^r(t)$ during trial r such that

$$E\{M^r(t)\} = \beta^r b + \rho^r p(t), \tag{2.11}$$

where β^r and ρ^r are scalar constants that depend on the trial number r. The constants $\{\beta^r\}$ will be used to track the trial-dependent changes in background firing rate seen in the experiment, while

---

[7] The "gain" measures how many spikes each neuron produces in response to the stimulus driving it.

[8] Note that since the PSTH is a *peri*-stimulus time histogram, there is a period of data collection that takes place before the start of the stimulus on each trial.

the constants $\{\rho^r\}$ will be used to track the trial-dependent changes in the gain of the stimulus-induced response. There are as many equations of the form of (2.11) as there were trials in the experiment. Notice that all that has been determined in the model is the expected value of $M^r(t)$. In what follows, no assumptions beyond this are necessary. In particular, it is not necessary for the model cell to be Poisson. Now, to be consistent with the number of pre-stimulus and total spikes observed in each trial of the experiment, estimate β^r and ρ^r for the model by setting them so that

$$n_b^r = \beta^r \sum_{t<t_0} b = \sum_{t<t_0} E\{M^r(t)\}. \tag{2.12}$$

$$n^r = \beta^r \sum_t b + \rho^r \sum_t p(t) = \sum_t E\{M^r(t)\}. \tag{2.13}$$

The first equation implies that $\langle \beta^r \rangle = 1$; the second implies that $\langle \rho^r \rangle = 1$. Thus, when averaged over trials, equation (2.11) is equal to the original PSTH:

$$\langle E\{M^r(t)\}\rangle_r = \langle \beta^r \rangle\, b + \langle \rho^r \rangle\, p(t) = b + p(t). \tag{2.14}$$

The description of the cell defined by b, $p(t)$, $\{\beta^r\}$, and $\{\rho^r\}$, is a form of "enhanced" PSTH that describes background and stimulus-dependent responses separately, and also describes the experimentally observed trial-dependent variations in the magnitudes of these two components. Interpreted as models, two such descriptions (one for each cell) can be used to predict the covariogram due to excitability correlations: Let the random variables in each model be independent of those in the other model. Then the two model cells can only interact if there are covariations in their β^r and ρ^r parameters. These in turn are determined by the experimental data, as described above. Adding the subscripts 1 and 2 to differentiate between the two cells, dropping the superscripts r for legibility and following equations (2.11) (2.10) and (2.2), the covariogram predicted by the models, which we will call V_e for "excitability covariogram," is

$$\begin{aligned} V_e \;=\; & (\langle \beta_1\beta_2 \rangle - 1)\, b_1 \otimes b_2 \;+\; (\langle \beta_1\rho_2 \rangle - 1)\, b_1 \otimes p_2 \;+ \\ & (\langle \rho_1\beta_2 \rangle - 1)\, p_1 \otimes b_2 \;+\; (\langle \rho_1\rho_2 \rangle - 1)\, p_1 \otimes p_2. \end{aligned} \tag{2.15}$$

The two models satisfy

$$\underbrace{\sum_t E\{M_1^r(t)\}}_{\text{model}} = \underbrace{n_1^r}_{\text{experiment}} \qquad \text{and} \qquad \underbrace{\sum_t E\{M_2^r(t)\}}_{\text{model}} = \underbrace{n_2^r}_{\text{experiment}}, \tag{2.16}$$

so the expected DC value of V_e will exactly match the DC value of the covariogram of the experimental data.

Two examples of making such estimates, and the result of subtracting them from the original covariograms, are shown in Figure 2.5. The top row of the Figure shows this for spike trains made with excitability correlations only. After subtracting the excitability covariogram, all remaining deviations are within the significance limits. The bottom row shows the estimate, and the covariogram after its subtraction, for spike trains constructed with both excitability and spike timing correlations. Since the rasters were artificial, all parameters used in their construction are known. Each of the two components of the covariogram can therefore be predicted (based on knowledge of the methods used to construct the rasters) and compared to the data-based estimates. Only the spike timing component prediction and estimate are directly compared in the Figure (bottom right). The grey line is the covariogram predicted based on the spike timing parameters alone. It matches the residual covariogram well: subtracting the excitability correlations estimate has accurately revealed the spike timing component of the covariogram.

In experimental data, excitability correlations may be mixed with spike timing ones, and confound the measurement of the latter (Das and Gilbert, 1995). The procedure described in this section can be used to unconfound the measurement.

2.4.2 Identifying the presence of latency correlations

How can we identify whether a peak in a covariogram was caused by latency correlations? Let the two sets of spike trains involved be $\{S_1^r(t)\}$ and $\{S_2^r(t)\}$. Assume there is a significant peak in their covariogram. If the peak is due only to joint time-shifts in the spike trains (*latency correlations*), then by definition (see section 2.3.2 and appendix 2.8) the spike trains can be described as

$$S_1^r(t) = U_1^r(t + \Delta t^r) \qquad \text{and} \qquad S_2^r(t) = U_2^r(t + \Delta t^r) \tag{2.17}$$

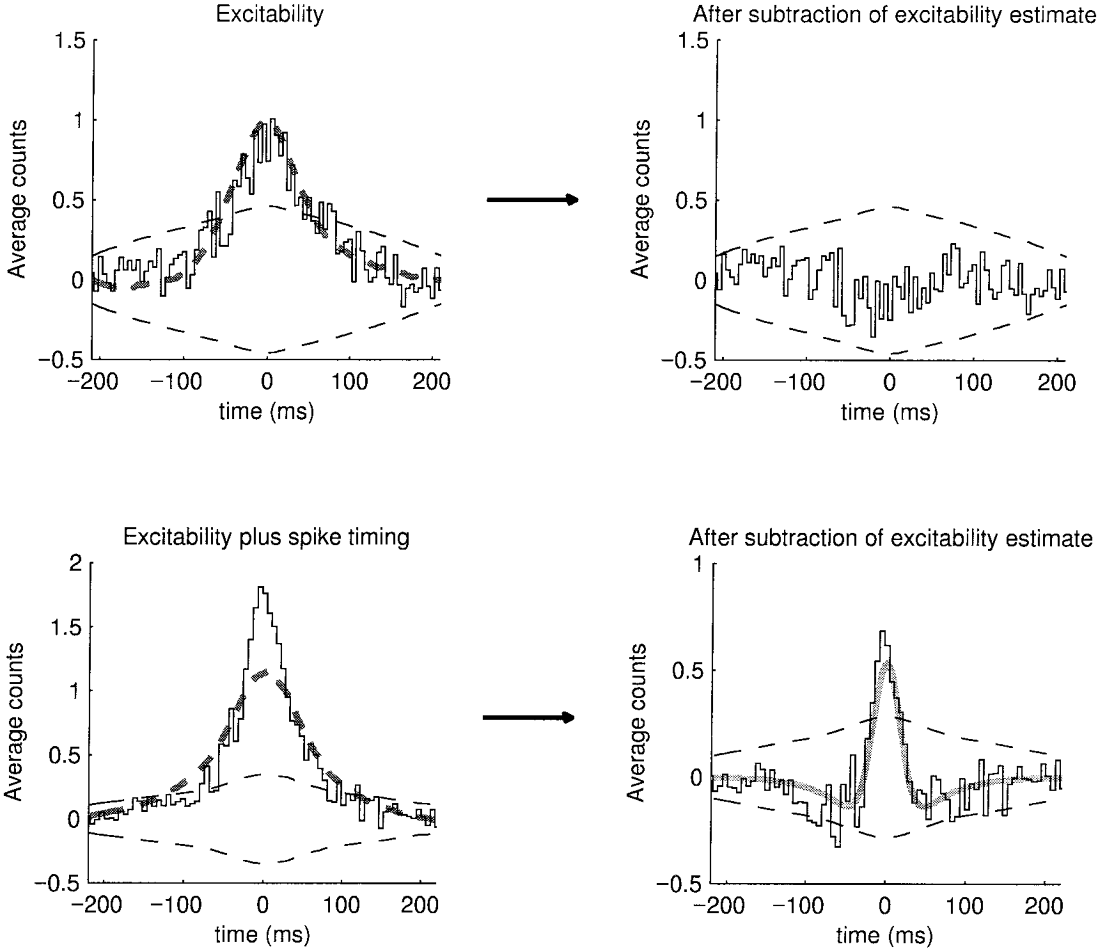

Figure 2.5: Estimating excitability correlations from the spike train data. **Top row:** On the top left is the same covariogram as shown in figure 2.2, an example of pure excitability correlations; overlaid on it as a thick dashed line is the excitability covariogram estimate V_e obtained using equation (2.15). Unlike the solid grey line in figure 2.2, the thick dashed line here was calculated without any knowledge of how the spike trains were generated. The top right shows the same covariogram after subtraction of the estimated excitability covariogram. All remaining deviations from zero are within the significance limits. **Bottom row:** Mixed-type correlations. On the bottom left is the covariogram of rasters constructed with both spike timing *and* excitability correlations. The shape does not obviously indicate two types of correlations; however, we can use more than just the shape to separate the two. The thick dashed line is the excitability covariogram estimate calculated using equation (2.15). On the bottom right this estimate has been subtracted, leaving behind a clear peak, indicative of a correlation other than an excitability correlation. We know here that this is a spike timing correlation, and can predict the expected covariogram shape based on knowledge of the spike timing correlation parameters used to construct the rasters. The predicted shape is shown as a thick grey line. It matches the residual covariogram well. Subtracting the excitability covariogram estimate has accurately revealed the spike timing component of the covariogram.

where $\{U_1^r(t)\}$ and $\{U_2^r(t)\}$ are spike trains drawn from two independent distributions and $\{\Delta t^r\}$ is some set of time shifts. In particular, the covariogram of U_1 and U_2 must be, within sampling noise, zero for all t.

The preceding observation suggests using the following method: search for time shifts $\{\Delta t^r\}$ such that deviations from zero (i.e. peaks or troughs) in the covariogram of

$$U_1^r(t) = S_1^r(t - \Delta t^r) \qquad \text{and} \qquad U_2^r(t) = S_2^r(t - \Delta t^r) \tag{2.18}$$

are minimized. Let the set of time shifts which minimizes the peaks and troughs be indicated by $\{\Delta \hat{t}^r\}$. If after the minimization the deviations are still statistically significant, the original peak in the covariogram of $\{S_1^r(t)\}$ and $\{S_2^r(t)\}$ cannot be explained in terms of joint latency shifts alone; but if the deviations are all statistically *in*significant, the result suggests that it can.[9] Furthermore, the result indicates that the PSTHs of the two independent, unshifted spike trains are $\langle U_1^r(t)\rangle = \langle S_1^r(t - \Delta \hat{t}^r)\rangle$ and $\langle U_2^r(t)\rangle = \langle S_2^r(t - \Delta \hat{t}^r)\rangle$, and the time shifts which led to the latency peak are $\{\Delta \hat{t}^r\}$. Knowledge of the PSTHs and time shifts is enough to predict the covariogram (see appendix 2.8). This prediction can then be compared to the original covariogram.[10,11]

The results of such a search run on the rasters of Figure 2.3 are shown in Figure 2.6. The search for time shifts was done by minimizing the scalar function $F = \int V^2(t)\,\mathrm{d}t$, which is a minimum when $V(t) = 0$ for all t. Here $V(t)$ is the covariogram of U_1 and U_2. This was not the only possible choice of $F(\cdot)$. Figure 2.6 shows the rasters $U_1(t)$ and $U_2(t)$ and their covariogram, demonstrating that it is, within sampling noise, zero for all t. Similarly, Figure 2.6 shows that the normalised JPSTH of $U_1(t)$ and $U_2(t)$ has no significant structure different from zero. This JPSTH should be compared to Figure 2.3. The Figure also shows the covariogram prediction based on $\langle U_1^r(t)\rangle$, $\langle U_2^r(t)\rangle$, and $\{\Delta \hat{t}^r\}$, overlaid on the original covariogram.[12]

---

[9] If there are excitability as well as latency correlations in the data, the search must evaluate and minimize deviations in the covariogram that remains after subtraction of the excitability estimate (section 2.4.1), calculated anew for each set of times $\{\Delta t^r\}$.

[10] In the minimization and in the prediction Δt^r is the same for both cells in each trial r. Constraining Δt^r to be the same for both cells is equivalent to the assumption that the peak seen in the original covariogram is most easily reproduced by latency correlations when the time shifts are the same for both cells. The rationale behind this is that latency correlation peaks are biggest when the time shifts for the two cells are most correlated, i.e. when the time shifts for both cells are the same.

[11] If reliable estimates for the latency shifts Δt^r for each cell are available, these should be used in place of of $\Delta \hat{t}^r$: the search generalizes the method to the case where such estimates are not available or are unreliable.

[12] Details of the search method, while important in practice, are not explored in this paper. The method may or may not use measures of latency specific to the problem in question (e.g. one might to choose to start the search by aligning all the bursts in the record, or, as suggested by Peter König (personal communication) aligning the rasters so the width

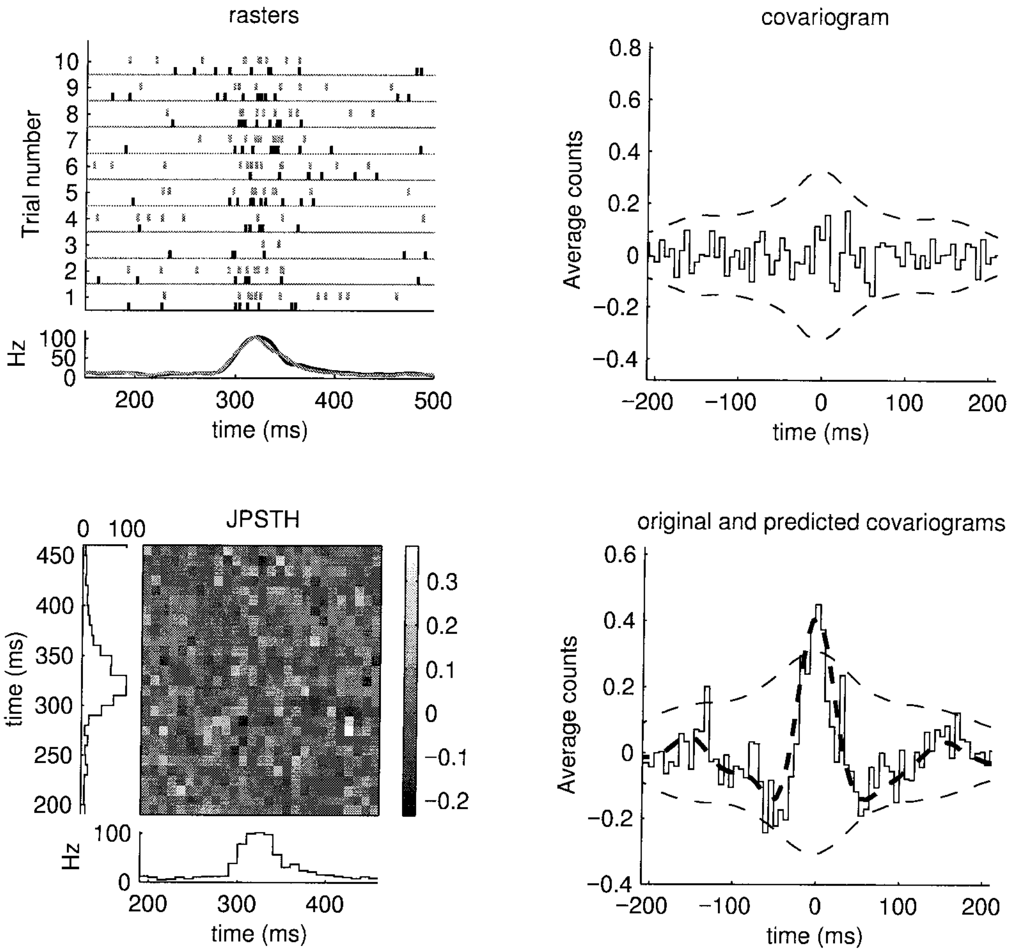

Figure 2.6: Latency search results. The top left shows the same rasters as in Figure 2.3, after shifting by the times found by the latency search described in the text. For example, the fifth raster has been shifted 30 ms to the left. The top right shows the cross-covariogram after the latency search. All deviations from zero are within the significance limits. The bottom left shows the JPSTH after the latency search. There is no structure significantly different from zero left (compare to Figure 2.3). The bottom right shows the original cross-covariogram from Figure 2.3 (thin solid line, binned) together with the covariogram predicted by the hypothesis found by the latency search method described in the text (thick dashed line).

Showing that a covariogram is well explained by latency correlations is not equivalent to showing that it was caused by them. It merely establishes that the latency correlations hypothesis is not one that may be ignored.

Not every covariogram can be explained in terms of latency correlations. Figure 2.7 shows an example using spike trains constructed using spike timing correlations. There is no set of latency shifts that fully accounts for the peak.

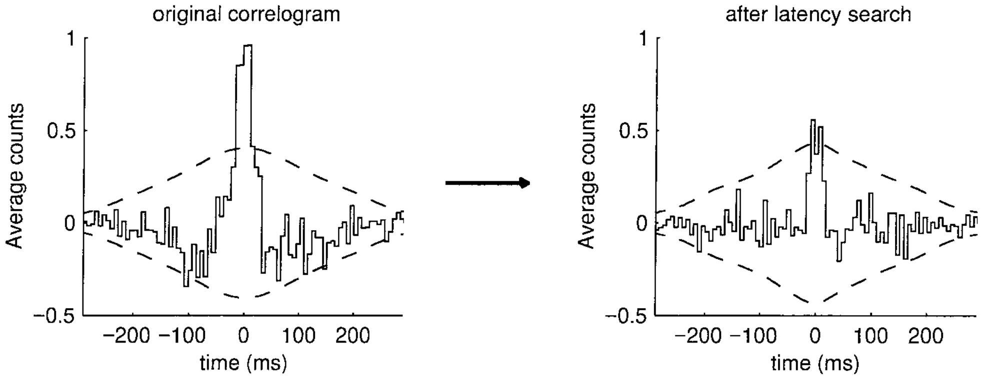

Figure 2.7: Covariogram of rasters built with spike timing correlations and result of the latency search procedure. There is no set of latency shifts that can fully account for the peak in the original covariogram. Note that the procedure is biased: because it searches for a minimum, it is guaranteed to reduce the peak by some amount. One might interpret the result shown here as indicating that the peak in the original covariogram is due to a mix of spike timing and latency correlations, but all that has been shown is that the covariogram alone does not distinguish between the hypothesis that the peak is due to spike timing correlations and the hypothesis that the peak is due to a mix of spike timing and latency correlations. Parsimony might suggest spike timing correlations only.

2.5 Discussion

When the firing patterns of two neurons covary, their shuffle-corrected correlogram (covariogram) will reflect this by not being zero for all t. Peaks or troughs in the covariogram can be interpreted in terms of individual spike-to-spike interactions between the two cells (*spike timing correlations*). Such interactions are complex in the sense that during every trial, information about many parameters (the firing times of many individual spikes) is shared between the two cells. However, there are other more global parameters, which describe the overall envelope of the response of

---

of the individual PSTHs is minimized); here, F was minimized through a simple iterated line search along coordinate directions.

each cell, about which information may be shared. Covariations between those parameters will also be reflected in peaks or troughs in the covariogram. I have presented two examples in this paper: covariations in an overall response timing parameter (*latency correlations*) and covariations in an overall gain parameter (*excitability correlations*).[13] These were chosen both because of their simplicity and because of their neurophysiological relevance. The covariograms generated by covariations in latency or excitability can be strikingly similar to covariograms generated by spike-to-spike interactions. This implies that one must be cautious when interpreting covariograms, since the three types of correlations can have similar covariograms yet may be caused by very different mechanisms. The three different types of correlations also have very different effects on postsynaptic cells receiving input from cells correlated in one of the three ways (C.D. Brody, manuscript in preparation). Procedures to distinguish between the three types of correlations were presented in the paper.

In the limit of one spike per trial, latency and spike timing correlations are identical. These two types of correlations, in particular, are thus better understood as two ends of a common spectrum. They both depend on sharing information about timing parameters: but while latency correlations share information about only one timing parameter per trial, spike timing correlations can share information about many more. Covariograms of two synchronously oscillating neurons (Singer and Gray, 1995a), where there is also a single timing parameter (oscillation phase) about which information is being shared on each trial, can also be seen as a form of latency correlations.

What kind of covariogram shapes may we expect from latency and excitability correlations? Since they both operate on the overall envelopes of the response (the PSTHs), the shapes of their covariograms are closely tied to these. The temporal structure in the covariograms will have the same time scale as the temporal structure in the PSTHs: thin, sharp peaks will be seen in latency or excitability correlograms only if the underlying PSTHs also have thin, sharp peaks. Covariogram peaks with a width on the order of ten milliseconds or less (often considered as being caused by ionotropic synaptic connections, e.g. Alonso et al. 1996) that are based on recordings with PSTH widths on the order of many hundreds of milliseconds or more are unlikely to have been caused by latency or excitability correlations. Broader peaks, on the other hand, may well have been caused by such interactions. It is important to note here that latency time shifts blur out the sharpness

---

[13] A passing reference to excitability correlations, as a possible confounder of spike timing correlation measurements, was made in Aertsen et al. (1989), This passing reference is cited further below. Perkel et al. (1967) also mention excitability correlations, but only in the context of unstimulated, spontaneously firing neurons; they concluded that in that regime excitability correlations are unimportant. To my knowledge, latency correlations have not appeared anywhere in the literature.

of the relevant peaks in the PSTHs. This is especially true for bursty neurons, which may have very sharp onset responses which are nevertheless not precisely time-locked to the stimulus. In such cases, thin sharp peaks may well be caused by latency covariations even though the two measured PSTHs need not be sharply peaked. This is because it is the "underlying" PSTHs, not yet smoothed by the latency variations (e.g. $\bar{U}_1$ and $\bar{U}_2$ of appendix 2.8), which must be sharply peaked in order to see a correspondingly sharp peak in the correlogram.

What physiological mechanisms could cause latency or excitability correlations? Conceivable mechanisms range from ones unrelated to the stimulus (e.g. slow drifts in the general physiological state of the experimental animal, or, for visual neurons, uncontrolled drifts in eye position (Rodieck et al., 1967; O'Keefe and Berkley, 1991)) to those directly related to processing of the stimulus (e.g. adaptive processes specific to the neurons being studied, an example of which would be adaptation at the synapses driving the neurons being recorded from, or correlated changes in the response mode of the neurons). In either case, detecting the presence of the correlations is a necessary first step in identifying the mechanisms causing them.

Section 2.3.1 pointed out that if all trials are of the same length of time, *the DC value of a shuffle-corrected correlogram is proportional to the covariation, over different trials, of the mean firing rate of the two cells.* Large, positive, DC values indicate strong covariations; negative DC values indicate anti-covariations. The same is true of shift-corrected correlograms. A brief glance at a shuffle- or shift-corrected correlogram allows one to tell whether the mean firing rates covaried significantly. Many shuffle- or shift-corrected correlograms found in the neurophysiological literature show such large positive DC values, indicating that covariations in the mean firing rates were indeed present (e.g. Kruger and Aiple 1988; Alloway et al. 1993; Hata et al. 1993; Ghose et al. 1994; Sillito et al. 1994; Nowak et al. 1995; Munk et al. 1995). The DC value of the corrected JPSTH is also proportional to the covariation over different trials of the mean firing rates of the two cells, which led Aertsen et al. (1989) to observe that "For a linear trend in *both* neuron firing rates [...] the PST-predictor can either be an underestimate or an overestimate, depending on whether the trends in the two neurons go in the same or in obverse directions" (page 914). The PST predictor in the JPSTH is analogous to the shuffle corrector K used in covariograms. If such joint trends in the firing rates exist, subtracting the PST predictor from the raw JPSTH (equation (2.5)) can leave behind a peak (or a trough). However, it would be inappropriate to consider that a failure of the PST predictor: the peak or trough is quite correctly signaling the fact that the firing of the two neurons is not independent. In that sense the PST prediction/ shuffle correction has performed exactly as desired. What would be incorrect would be to interpret the peak as a spike timing correlation, since it is a peak due to an excitability correlation.

Latency and excitability are not the only global parameters that could covary in the two cells. They are two of the simplest and most general. The most obvious remaining simple case, not treated in this paper, is that of covarying changes in the time *length* of the response of the two cells. A quantitative method for detecting whether the presence of covarying changes in the time length of the response can account for observed peaks in a correlogram remains to be fully developed. Together with the cases treated in this paper, the kinds of simple "warps" of the PSTH response envelopes that would then be covered are translations, both in time (x-axis; latency correlations) and in firing rate (y-axis; excitability correlations in background firing rates); and dilations/compressions, again both in time (time length correlations) and in firing rate (excitability correlations in stimulus-induced responses). The various warps may occur together in different combinations, and even in different correlated combinations. An example is the following hypothetical scenario: suppose the membrane resting potential of two cells covaries over different trials. If the two cells are well described by leaky integrate-and-fire models, trials during high resting potentials will result both in short latencies (the driving force can quickly drive the neurons to spike, since the resting potential is closer spiking threshold than usual) and in high excitabilities (the driving force will lead to more spikes than usual), whereas trials during low resting potentials will result in long latencies paired with low excitabilities. In this scenario, peaks in the correlogram of the two cells would be well accounted for by trial-dependent, paired variations in both latency and excitability. This would be found by the methods described in sections 2.4.1 and 2.4.2. The two types of variations being paired (that is, that the latency would covary with the excitability) would then suggest the existence of a *single* parameter (resting membrane potential) underlying both the latency and the excitability correlations– a possibility that could be further explored through biophysically realistic models that explicitly include a description of resting membrane potential. What this hypothetical scenario illustrates is that the basic warps described in the paper are not to be taken as an end product. Rather, they are to be taken as flexible descriptions of covariations of the global response properties of two (or more) neurons. If a particular experimental correlogram is well described in these terms, then that description should be examined further, to be used as an aid to understanding the interactions between the neurons in terms of biophysically and/or computationally relevant parameters and processes.

While translations and dilations can describe many simple changes in the global response properties of the neurons, there are changes in global response properties that will not be well described in those terms. For example, one might be studying two neurons whose responses change together from trials where the cells have an initial burst onset in response to the stimulus, to trials where neither cell bursts. Yet each neuron might, for both cases, fire similar total numbers of spikes, at similar latencies, and in similar overall lengths of time. Thus the variations would not be well described by translations and dilations of the PSTHs. But since the firing patterns of the two cells would be covarying, this case, too, would produce a peak in the shuffle-corrected correlogram of the two cells. One must carefully examine the rasters of the responses to look for such possibilities. Even though, as a second-order statistic, a correlogram is always constructed by considering two spikes at a time, peaks or troughs in it are *not* necessarily caused by individual spike-to-spike interactions.

Acknowledgements

I am grateful to John Hopfield and members of his group at Caltech, Sanjoy Mahajan, Sam Roweis and Erik Winfree, for discussion and critical readings of the manuscript. I also thank John Hopfield for his support and for pointing out that paired latency and excitability correlations could be due to covarying resting membrane potentials. I thank George Gerstein, Kyle Kirkland and Adam Sillito for discussion. This work was supported by a Fulbright/CONACYT graduate fellowship and by NSF Cooperative Agreement EEC-9402726.

All simulations and analyses were done in MATLAB 5 (Mathworks Inc, Natick, MA). The code for all of these, including in particular the code to reproduce each of the figures, can be found at `http://www.cns.caltech.edu/~carlos/correlations.html`.

2.6 Appendix

2.7 Generating spike trains

All the methods used in this paper to generate artificial spike trains relied on generating trains of spikes from Poisson units with time-varying firing rates. Given a Poisson unit with a firing rate $\lambda(t)$, the probability that it will produce a spike in the infinitesimal time interval $[t, t+\delta t]$ is defined as $\lambda(t)\delta t$. Such units were approximated by finely discretizing time ($\Delta t = 0.1$ ms) and then placing a spike at each time-step with probability $\lambda(t) \cdot \Delta t$. For the rates used here, $\lambda(t) \cdot \Delta t$ was always less than 0.02 (maximum firing rate 200 Hz).

2.7.1 Poisson model of excitability correlations

To generate a set of rasters with excitability correlations, the firing rates of two simulated Poisson units were set on each trial to be $\gamma^r \lambda(t) + b$, where γ^r was a scalar constant that varied with trial number r. On each trial the trains for the two units, though based on the same time-varying firing rate, were generated independently. Under those conditions, the expected PSTH of each cell is $E\{S(t)\} = \langle \gamma^r \rangle \, \lambda(t) + b$ and the expected covariogram is

$$E\{V\} = (\langle \gamma^2 \rangle - \langle \gamma \rangle^2) \cdot \lambda \otimes \lambda \tag{2.19}$$

For Figure 2.2, two hundred trials were generated, γ^r was drawn anew on each trial from a Gaussian distribution with mean 1 and standard deviation 1 (γ^r was set to zero on each trial if it was negative), b was 35 Hz, and $\lambda(t)$ was an alpha function that began at 70 ms and peaked at 100 ms with a value of 70 Hz: $\lambda(t) = (70 \text{ Hz}) \frac{t-70}{30} \exp(\frac{100-t}{30})$ if $t > 70$, zero otherwise, with t in milliseconds.

2.7.2 Poisson model of latency correlations

Rasters with latency correlations were generated by setting the firing rates of two independent cells to be $\lambda(t)$ on every trial, and then shifting both spike trains in each trial by an amount of time t^r that was drawn anew for each trial from a distribution $\Pi(t)$. In addition, an independent unshifted spike train generated with constant firing rate b was added to each cell in each trial, to represent random background firing. Under those conditions, the expected PSTH for each cell is $E\{S\} = \lambda \otimes \Pi + b$, and the expected covariogram is

$$E\{V\} = (\lambda \otimes \lambda) \otimes (\delta(0) - \Pi * \Pi) \tag{2.20}$$

where * indicates convolution and $\delta()$ is the Dirac delta function or the Kronecker delta δ_{0t}, according to whether the calculation is being done in continuous or discrete time.

For Figure 2.3, 200 rasters were generated, Π was a Gaussian with mean zero, standard deviation 15 ms, b was 10 Hz, and $\lambda(t)$ was $(100 \text{ Hz}) \exp(\frac{(t-100)^2}{2 \cdot 40^2})$ for $t > 100$, zero otherwise, with t in milliseconds.

2.7.3 Poisson model of spike timing correlations

Each trial of rasters with spike timing correlations was generated by first generating a spike train from a unit with firing rate set to be $\lambda(t)$. This spike train then had its spike times jittered twice, once to produce the spikes for cell 1 and once again to produce the spikes for cell 2. In addition, each cell had an independent spike train generated with constant firing rate b added to its spike train, to represent random background firing. Jittering was done by adding a random time, drawn i.i.d. for each spike from a distribution $j(t)$, to the time of each spike in the train. The expected PSTH for each cell is then $E\{S\} = j \otimes \lambda + b$. Let the expected number of spikes in each trial that have a corresponding spike in the other cell's spike train be labelled $N_s = \sum_t \lambda(t)$. The expected covariogram is

$$E\{V\} = N_s[j \otimes j] \tag{2.21}$$

For Figure 2.4, j was a Gaussian with mean zero and standard deviation 12 ms, b was 10 Hz, and $\lambda(t)$ was a Gaussian-shaped function with mean 100 ms, standard deviation 30 ms, and a peak value of 70 Hz.

Spike timing without excitability correlations

The above procedure for spike timing correlations includes some excitability correlations: the total number of spikes in each trial covaries in the two cells, since for any given trial both trains contain jittered versions of the same train drawn from $\lambda(t)$. A different method that does not create excitability correlations was used for some of the simulations. When drawing the unjittered spike train based on the time-varying rate $\lambda(t)$, the spike train was *constrained* to have exactly N_s spikes on every trial. This reduces the expected covariation in the number of spikes for the two cells to zero. The expected covariogram is now

$$E\{V\} = N_s[j \otimes j] - \frac{1}{N_s}\lambda \otimes \lambda \tag{2.22}$$

Note that since $\sum_t \lambda(t) = N_s$ and $\sum_t j(t) = 1$, $\sum_t E\{V\} = 0$.

2.8 Latency Correlations

This appendix presents mathematical details regarding latency correlations.

Let the responses of the two neurons for each trial r be described by

$$\begin{aligned} S_1^r(t) &= U_1^r(t + t_1^r) \\ S_2^r(t) &= U_2^r(t + t_2^r), \end{aligned}$$

where $U_1(t)$ and $U_2(t)$ are independent point processes, the times t_1^r and t_2^r are drawn from a probability distribution $P(t_1, t_2)$, and all three of U_1, U_2 and P are stationary over different trials. The overall process is stationary with respect to trial number. The expected value over many trials of $U_1(t)$ and $U_2(t)$ will be written as $\bar{U}_1(t)$ and $\bar{U}_2(t)$ and can be thought of as PSTHs.

We wish to determine the expected covariogram $E\{V\}$ of S_1 and S_2. Since the latency time shifts need not be integer multiples of any particular bin size used, the description below will be in terms of continuous time. The cross-correlogram for trial r in a given experiment is

$$C^r(t) = \int \mathrm{d}\tau \, U_1^r(t + \tau + t_1^r) U_2^r(\tau + t_2^r) \tag{2.23}$$

Therefore, since the process is stationary over trials,

$$\begin{aligned} E\{R\} &= E\{\langle C^r(t)\rangle\} = E\{C^r(t)\} \\ &= \int \mathrm{d}\tau \, \mathrm{d}t_1^r \, \mathrm{d}t_2^r \; E\{U_1^r(t + \tau + t_1^r) U_2^r(\tau + t_2^r)\} P(t_1^r, t_2^r) \\ &= \int \mathrm{d}\tau \, \mathrm{d}t_1 \, \mathrm{d}t_2 \, \bar{U}_1(t + \tau + t_1) \bar{U}_2(\tau + t_2) P(t_1, t_2) \end{aligned} \tag{2.24}$$

A change of variables leads to writing the expected raw cross-correlogram as

$$E\{R\} = \int \mathrm{d}t_1 \, \mathrm{d}t_2 \; P(t_1, t_2) \; C^U(t + t_1 - t_2), \tag{2.25}$$

where

$$\begin{aligned} C^U(t) &= \int \mathrm{d}\tau \, \bar{U}_1(t + \tau) \bar{U}_2(\tau) = \bar{U}_1 \otimes \bar{U}_2 \\ &\equiv \text{expected cross-correlogram of } U_1 \text{ and } U_2. \end{aligned} \tag{2.26}$$

The two expected PSTHs are

$$
\begin{aligned}
E\{S_1(t)\} &= \int \mathrm{d}\tau\, \bar{U}_1(t+\tau)P_1(\tau) = \bar{U}_1 \otimes P_1 && (2.27)\\
&\text{and}\\
E\{S_1(t)\} &= \int \mathrm{d}\tau\, \bar{U}_2(t+\tau)P_2(\tau) = \bar{U}_2 \otimes P_2, && (2.28)
\end{aligned}
$$

where $P_1()$ and $P_2()$ are the marginal distributions of t_1 and t_2, respectively. The expected shuffle corrector $E\{K\}$ can be shown to be

$$E\{K\} \;=\; E\{S_1\} \otimes E\{S_2\} + \frac{1}{N}[E\{R\} - E\{S_1\} \otimes E\{S_2\}], \qquad (2.29)$$

where N is the number of trials in each repetition of the experiment. The expected cross-covariogram is then

$$
\begin{aligned}
E\{V\} &= \frac{N-1}{N} \cdot [E\{R\} - E\{S_1\} \otimes E\{S_2\}]\\
&= \boxed{\frac{N-1}{N} \cdot [E\{R\} - (U_1 \otimes P_1) \otimes (U_2 \otimes P_2)]} && (2.30)
\end{aligned}
$$

where the $E\{R\}$ is described by equation (2.25). For large N the constant prefactor $(N-1)/N$ is negligibly close to one.

There are three simple limiting cases:

Case (1): $P(t_1, t_2) = \Pi(t_1)\delta(t_2 - t_1)$, where $\delta()$ is the Dirac δ function. That is, the times t_1 and t_2 are always the same. This is the case described ins ection 2.3.2 and used in the simulation of Figure 2.3. From equation (2.25),

$$E\{R\} = \int \mathrm{d}\tau\, \Pi(\tau)C^U(t) = C^U(t). \qquad (2.31)$$

Assuming $N \gg 1$, equation (2.29) gives

$$E\{K\} \approx E\{S_1\} \otimes E\{S_2\} = \bar{U}_1 \otimes \bar{U}_2 \otimes (\Pi * \Pi) = C^U \otimes (\Pi * \Pi), \qquad (2.32)$$

where $\Pi * \Pi$ indicates the convolution of Π with itself. Sharp peaks or troughs in the raw cross-correlogram are smoothed away in the shuffle corrector to remain outstanding in the expected covariogram $E\{V\}$, as was described in section 2.3.2.

Case (2): $P(t_1, t_2) = P_1(t_1)P_2(t_2)$, that is, the two times are independent. Equation (2.24) leads to $E\{R\} = E\{S_1\} \otimes E\{S_2\}$, which implies that $E\{V\} = 0$. This is exactly what we would expect, since if t_1 and t_2 are independent, then the two cells are in fact completely independent.

Case (3): $P(t_1, t_2) = \Pi(t_1)\delta(t_2 + t_1)$, that is, $t_1^r = -t_2^r$ so the latencies of the two cells are perfectly *anti*-correlated. This is the inverse of case (1), and a similar analysis to that of case (1) shows that the dominating feature in the covariogram is a trough indicative of the anticorrelation between the two cells.

Chapter 3 Spike train correlations in the LGN of the cat: Analysis

3.1 Summary of the chapter

Sillito, Jones, Gerstein, and West recently obtained spike train data from pairs (sometimes triads) of simultaneously recorded neurons in the dorsal lateral geniculate nucleus (LGN) of anesthetized cats. This chapter describes an analysis of that data set, with an emphasis on the interactions between pairs of neurons. The analysis done here extends that published in a previous report (Sillito, Jones, Gerstein, and West, 1994); the new analysis suggests an interpretation of the data significantly different to that published in that report. The differences are summarized in this chapter introduction (3.1) and in the Discussion chapter (5).

Section 3.2 will begin the chapter by describing some basic properties of single-cell responses in the LGN, and the experimental protocol used by Sillito et al. Subsection 3.2.1 is primarily intended for the reader unfamiliar with cat LGN and may be skipped by readers with more expertise. Subsection 3.2.2 then describes the experimental methods and protocols used by Sillito et. al. In very brief: Three separate electrodes, fixed in position with respect to each other, were used to record extracellularly from pairs (sometimes triads) of neurons in the LGN of anesthetized cats. The electrode tips formed a roughly equilateral triangle. The distance between the electrodes was chosen so that the receptive fields of the cells recorded from each electrode did not overlap with the receptive fields of cells recorded from other electrodes; receptive fields centers lay between 1° and 4° degrees apart from each other. Spike times of single neurons were isolated from each electrode's analogue signal using window discriminators. Once receptive field locations and types (On/Off/X/Y) were determined, the cells were driven by visual stimulation. This was most often a series of drifting sine wave gratings of various orientations, with trials at different orientations pseudo-randomly interleaved with each other. Interactions between pairs of cells were then assessed using covariograms and JPSTHs (quantified as in section 2.2).

The main result obtained by Sillito et al. was the discovery of significant positive peaks in the covariograms of spike trains of cells recorded from different electrodes, that is, of cells with separate, non-overlapping receptive fields.[1] A significant peak in a covariogram indicates that at

---

[1] Shift predictors and shuffle corrector methods (Perkel et al., 1967; Palm et al., 1988; Aertsen et al., 1989) were

least one parameter controlling the firing of cell 1 (call this parameter π_1) covaried with at least one parameter controlling the firing of cell 2 (call that parameter π_2). The time scale over which the parameters covaried will be reflected in the width of the peak in the covariogram. But if the spike train data is broken into different trials (as it is here), and the time scale over which π_1 and π_2 correlate is longer than the length of each trial, the width of the peak in the covariogram can no longer directly reflect the correlation's time scale: the complete covariogram itself has a total width of twice the trial length, so any peak in it is per force constrained to be at most as wide as that. Under these circumstances, then, the width of the peak in the covariogram may not accurately reflect the time scale of the correlation. Furthermore, if the neurons are driven periodically, the covariogram peak will *also* be modulated by that periodicity: during times when one of the two neurons is not producing spikes, the covariogram is constrained to be zero, regardless of whether the internal parameters π_1 and π_2 are correlated at those times. This again will lead to a covariogram peak with a width that does not directly reflect the time scale of the correlations. Section 3.3 describes how these considerations are taken into account in the analysis below.

Section 3.4 applies the ideas of section 3.3 to the experimental records in which covariogram peaks were found. The analysis is limited to those experiments in which drifting sine wave gratings were used as a stimulus, and is dedicated to determining the time scale and stimulus-orientation dependence of the correlations. The results show that (1) the time scale of the main component of the correlations is on the order of *tens of seconds to minutes*; (2) the correlations depend weakly, if at all, on the orientation of the stimulus; and (3) responses to stimuli of different orientations are correlated with each other. These are major differences with the previous report (Sillito et al., 1994), where the correlations were interpreted to have a time scale on the order of tens of milliseconds and to be highly dependent on stimulus orientation. See section 3.4.1 for an explanation of why the two sets of results are so different.

Section 3.5 uses the techniques of chapter 2 to show that most of the correlations in the data are well described in terms of combinations of latency and excitability correlations. (Latency and excitability correlations are fully described in chapter 2.) Explaining the covariogram peaks without recourse to a spike timing interpretation is again a major difference with the previous report (Sillito et al., 1994). In particular, the latency/excitability interpretation does not open the possibility of these LGN neurons participating in "binding" (Singer and Gray, 1995b), as the spike timing interpretation did.

---

always used. See chapter 2 for a discussion and explanation of these methods.

Since the latency and excitability of a neuron are both affected by its resting membrane potential, this suggests that the parameters covarying in the two neurons are their respective resting potentials. This possibility is introduced in section 3.6 and is explored through modeling in the next chapter.

3.2 Single-cell responses and description of the experiment

3.2.1 Single-cell responses in the LGN

The lateral geniculate nucleus (LGN) is the principal gateway of information from the retina to cortex, and as such has been extensively studied (reviewed in Sherman and Koch 1986; Sillito 1992; McCormick 1992; Sherman and Guillery 1996). In order to set the context for the multiple-cell recording experiments to be described and analysed below, I will first describe some key features of single-cell responses in the LGN, focusing in particular on cells in the A laminae of the LGN. All experimental data used for examples here was collected by Adam Sillito, Helen Jones, and colleagues. I thank them for generously making it available to me.

Receptive field properties

Cells in the LGN receive strong, topographically organized, input from the retina, and much of their characteristic physiology is derived from that input. Like cells in the retina, cells in the LGN have small, circularly symmetric receptive fields with a center-surround organization, and can be classed into "On" or "Off" type cells. The left panel of Figure 3.1 shows schematic receptive field profiles of these two types of cells. "On" cells produce spikes in response to a spot of light flashed on in the center of their receptive field, but are inhibited by similar spots flashed on in their surround. In contrast, Off cells produce spikes when a spot of light in their center is turned *off*, and are *excited* by spots of light turned on in their surround. The two types of cells are thus complementary to each other. If the cell is an X type cell (see below for distinction between X and Y cells), the response of the cell to spatiotemporally complex stimuli can be predicted by the convolution of its receptive field profile with the stimulus, followed by half-wave rectification of the result[2] (see Dan et al. 1996 for a striking example of the accuracy of this method). because convolution is a linear operation, X cells are thought of as quasi-linear. *Rectification* occurs for both X and Y cells because the firing rate of a cell cannot go below zero: if inhibited strongly, cells simply fall silent. The rectification is illustrated in the right panel of Figure 3.1, which shows the response of an LGN cell when stimulated by a drifting sine wave grating.[3] Assuming the cell is an On cell, when the bright bars of the grating pass over the center of the receptive field, the cell spikes; when the dark bars pass over the center of the receptive field, the cell is inhibited and falls silent. Thus, while the temporal profile of the input is that of a mean-zero sine wave, the temporal profile of the (average) response is closer to a half-wave rectified sine wave[4] (PSTH at bottom of the right panel of Figure 3.1). Much of the experimental data to be analysed below was collected using drifting sine wave gratings as stimuli, producing responses similar to those exemplified in Figure 3.1.

In addition to the On/Off distinction, cells in the LGN may be further divided into at least three classes, on the basis of both anatomical and physiological criteria (reviewed in Sherman 1985). The physiological properties of each class are again derived from the properties of the retinal cells providing input to the LGN, which are divided into the same classes. In the cat, there exist X, Y, and W-type cells, of which X and Y are the predominant population in the A laminae of the LGN. X cells have relatively small cell bodies and slow axon conduction velocities; they also have small receptive fields, and when studied with sine wave grating stimuli, respond optimally to high spatial and low temporal frequencies. Y cells have large cell bodies and fast axon conduction velocities; they have large receptive fields, and respond optimally to low spatial and high temporal frequencies.[5] A key distinction between X and Y cells is that the latter are *not*

---

[2] More precisely, if the receptive field has spatiotemporal profile $F(x, y, t)$ and the stimulus has luminance $S(x, y, t)$, the firing rate is proportional to $f(t) \cdot H(f(t))$, where

$$f(t) = \sum_{\tau > 0, x, y} F(x, y, -\tau) S(x, y, t - \tau) - \theta,$$

$H(z)$ is zero for $z \leq 0$ and 1 otherwise, and θ is a constant. In Figure 3.1, only the spatial component of F is shown, for clarity.

[3] Along the axis orthogonal to its orientation, a moving, or "drifting", sine wave grating has a luminance profile $I(x, t) \propto \sin(\omega_x x + \omega_t t + \phi)$.

[4] If a cell has a high background firing rate, it can represent much of the negative-going part of the sine wave through a decrease in its firing. Under the experimental conditions used by Sillito et. al., most LGN cells recorded from had low background firing rates and did not have that dynamic range available to them.

[5] This is reminiscent of the division into magnocellular and parvocellular streams in the monkey. The third type of cell, W, may perhaps be analogous to koniocellular cells in the monkey (Casagrande, 1994). W cells are mostly found in the C layers of the cat LGN and will not be considered here. Unlike monkey magnocellular and parvocellular streams, X and Y cells in the cat LGN are not physically segregated into different layers.

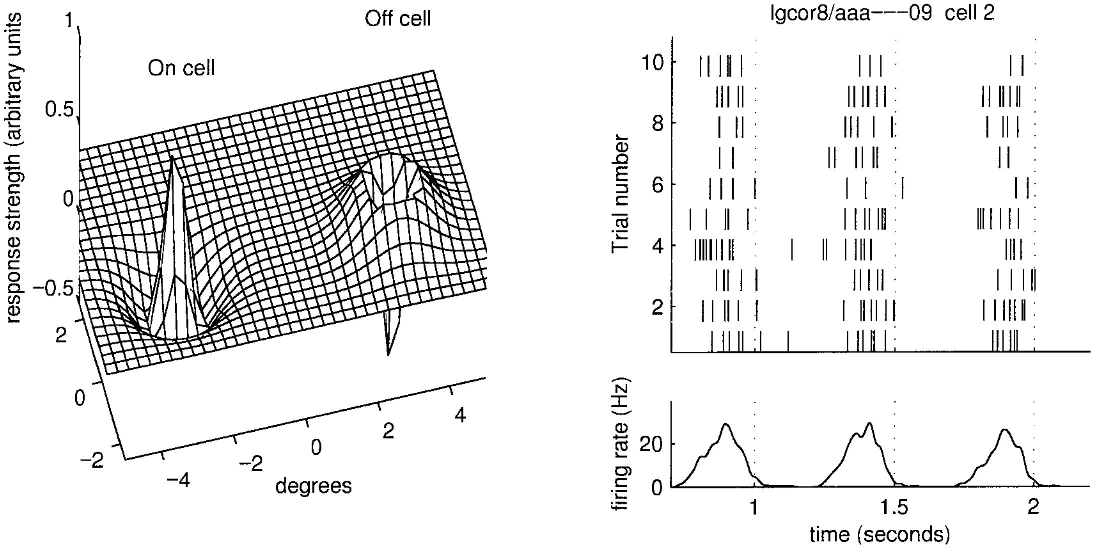

Figure 3.1: Typical LGN responses. The left panel shows a schematic plot of two receptive fields, one an On cell and the other an Off cell. x and y axes represent visual space; the z axis indicates the strength of the cell's response to a spot of light flashed on. Negative values on the z axis indicate inhibition. The right panel shows the responses of a real LGN neuron to stimulation by a drifting sine wave grating. Each vertical bar indicates a spike. 10 individual trials are shown, and below them the PSTH of 100 trials. The PSTH has been smoothed by a Gaussian with a standard deviation of 6 ms.

quasi-linear (Enroth-Cugell and Robson, 1966; Hochstein and Shapley, 1976). If stimulated with a stationary counterphasing grating,[6] an X cell's temporal response profile, being quasi-linear, will have most of its power at the temporal frequency of the stimulus (see upper left PSTH in Figure 3.2, obtained with a 1 Hz stimulus); in contrast, a Y cell's response will have significant power at the harmonics, in particular at twice the frequency of the stimulus (see lower left PSTH in Figure 3.2, obtained again with a 1 Hz stimulus). This is incompatible with the quasi-linear model.

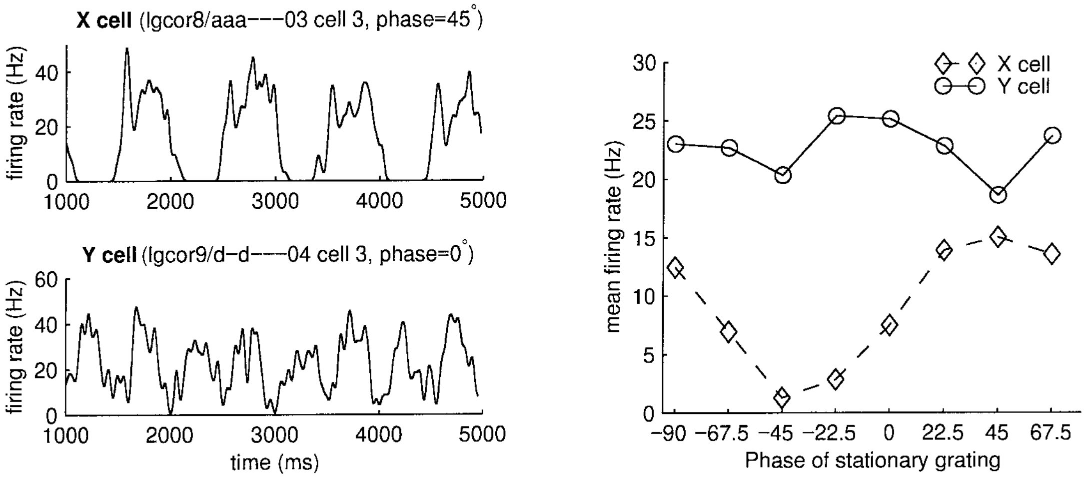

Figure 3.2: Differences between X and Y cells. The left panel shows two PSTHs, one for an X cell, the other for a Y cell, both stimulated by a 1 Hz stationary counterphasing grating. The X cell has most of its power at 1 Hz, but the Y cell shows significant power at the first harmonic, 2 Hz. The right panel shows data from the same two cells. The mean firing rate in response to a stationary counterphasing grating is plotted as a function of the spatial phase of the grating. The X cell's response is strongly affected by the phase of the grating, and goes through a "null point" where the cell responds with an almost-zero firing rate; the Y cell responds robustly at all phases.

In addition, the quasi-linear model predicts that if stimulated with a stationary counterphasing grating, the mean firing rate of the cell's responses will depend on the spatial phase of the grating. In particular, there will exist a spatial phase at which the driving force to the neuron will not be temporally modulated.[7] The phase at which the stimulus does not produce a temporal modulation

---

[6] Along the axis orthogonal to its orientation, a stationary counterphasing sine wave grating has a luminance profile $I(x,t) \propto \sin(\omega_x x + \phi_x) \cdot \sin(\omega_t t + \phi_t)$.

[7] If the receptive field of the X cell is spatiotemporally separable, then this is an inevitable consequence of the linear component of the model. Let $F(x,y)$ be the spatial receptive field profile of the cell; let $S(x,y;\phi)\sin(\omega t)$ be the visual stimulus, a stationary counterphasing grating at spatial phase ϕ; and let $F \cdot S(\phi) \equiv \sum_{x,y} F(x,y)S(x,y;\phi)$. Since S is a sine wave grating, $S(x,y;\phi_0+180) = -S(x,y;\phi_0)$ for all x,y and ϕ_0. Therefore $F \cdot S(\phi_0+180) = -F \cdot S(\phi_0)$. Since all functions here are continuous, there must be a value of ϕ, intermediate to ϕ_0 and $\phi_0 + 180$, for which

of the firing rate also leads to little or no elevation of firing rate above the spontaneous background rate of the cell. The right panel of Figure 3.2 shows a plot of mean firing rate of an X cell vs. spatial phase of a stationary counterphasing grating used to stimulate it. The phase at which the positive and negative components balance (the "null point") is close to 45 degrees. In contrast, a Y cell responds regardless of the spatial phase of the grating; an example of this is also shown in Figure 3.2. These two tests (power at twice the stimulus frequency, lack of a "null point") are often used to distinguish X and Y cells from each other. In the data analyzed below, both X and Y and both On and Off type cells were recorded from. As noted by Sillito et al. (1994), no cases of significant covariations between cells of different types (either X vs. Y or On vs. Off) were found.

Cellular properties

An important property of LGN cells that must be noted here (and that is shared with all other cells in the thalamus) is their capacity to produce bursts of Na^+ spikes, underlying which is a low-threshold Ca^{++} spike (Jahnsen and Llinas, 1984a,b). The low-threshold Ca^{++} channel responsible for these spikes can be well-modeled using Hodgkin-Huxley-like kinetics where the permeability of the channel is proportional to m^2h (McCormick and Huguenard, 1992; Huguenard and McCormick, 1992), m and h being the activation and inactivation gates of the channel, respectively.[8] Activation curves and time constants for these two gates, as described by McCormick and Huguenard (1992), are shown in Figure 3.3.

At normal resting potentials (-65 mV), the Ca^{++} channel is inactivated (see curve for h_∞) but can deinactivate if the membrane is hyperpolarized. (Due to the long time constant τ_h, a long hyperpolarization is required for the h gate to reach an appreciable value.) If the membrane potential has been hyperpolarized for long enough to deinactivate the channel, subsequent depolarization can activate the m gate (see curve for m_∞). Since the time constant for the m gate is faster than that of the h gate, after depolarizing the channel will be open during the time it takes h to reach h_∞, which can be tens to hundreds of milliseconds. During that time an influx of Ca^{++} occurs which further depolarizes the membrane in a (slow) all-or-none, voltage-spike-like manner. Eventually the depolarization leads to inactivation as h reaches a near-zero value (h_∞ is essentially

---

$F \cdot S(\phi) = 0$.

[8] m and h dynamics are given by

$$\dot{m} = (m_\infty(v) - m)/\tau_m(v)$$
$$\dot{h} = (h_\infty(v) - h)/\tau_h(v)$$

where v is membrane potential, and the functions $m_\infty, h_\infty, \tau_m, \tau_h$ are as shown in Figure 3.3.

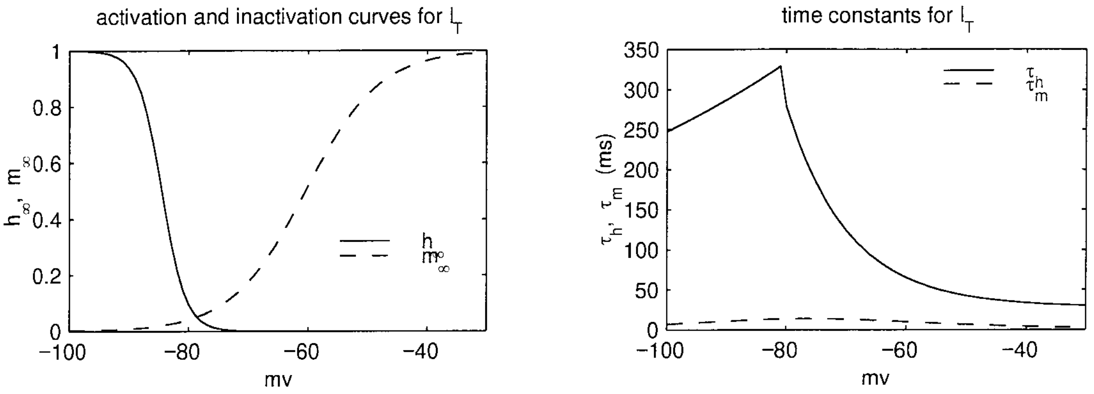

Figure 3.3: Activation and inactivation curves and time constants for a model of I_T, the low-threshold Ca^{++} channel current. The model is that of McCormick and Huguenard (1992).

zero for all voltages above -70 mV), and the Ca^{++} spike terminates. During the active part of the Ca^{++} spike, a burst of fast Na^{+} spikes riding on top of it may occur. Bursts, their modeling, and the mechanism of their generation will be treated further in the next chapter

The key features about this phenomenon to bear in mind for the present chapter are that (1) when activated, the Ca^{++} channel produces bursts of very high-frequency Na^{+} -spike firing in the range of 150 to 300 Hz and that these bursts last on the order of tens of milliseconds; and (2) control of activity or inactivity of the Ca^{++} channel rests with the cell's membrane resting potential (Jahnsen and Llinas, 1984a,b; Turner et al., 1994; Guido et al., 1992; Lu et al., 1992). Figure 3.4 shows an example of LGN spikes where some of these occur in brief bursts (highlighted by gray boxes). All rasters are from the same cell, illustrating that a single cell may exhibit both types of firing modes (known as *tonic* or *burst* firing). Bursts were identified as Ca^{++} channel-induced bursts using the extracellular criteria of Guido et al. (1992); Lu et al. (1992) (these criteria are: interspike interval before first spike in burst $\geq$ 100 ms, first interspike interval in burst $\leq$ 4 ms, subsequent interspike intervals $\leq$ 8 ms).

A change in resting membrane potential can lead to a change in the type of response, from tonic to bursty, or vice-versa. If two cells being recorded from change modes together, this is a type of covariation in their response, and will hence produce a peak in their covariogram (see chapter 2). One of the hypotheses that put forth in this thesis is that some of the covariogram peaks to be discussed below are due to precisely such paired changes in the two cell's firing modes.

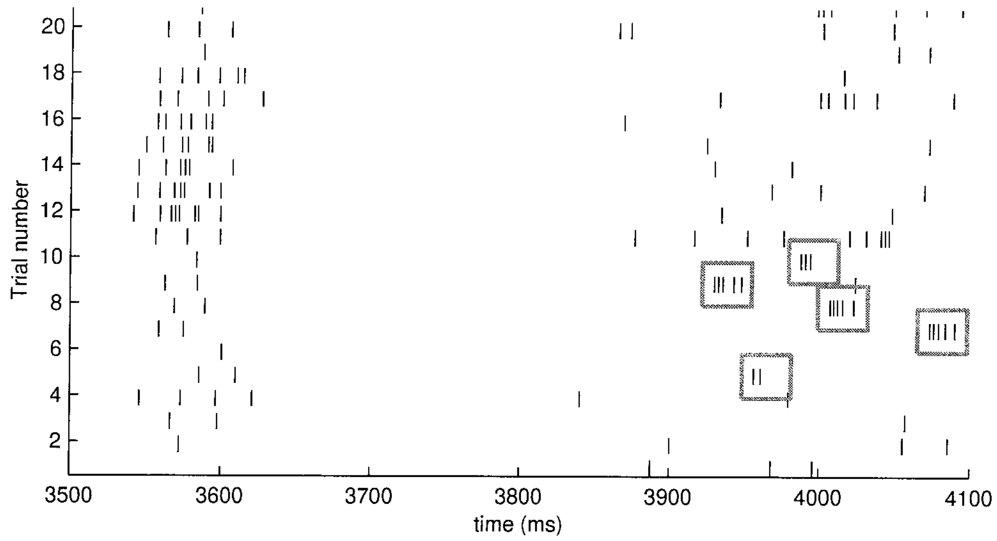

Figure 3.4: Tonic and burst firing. The same cells can produce tonic spikes and a few hundred ms later produce spikes in Ca^{++} bursts. The bursts have been surrounded by grey boxes. As in other figures in this section, the experimental data was provided by Sillito et al. and is taken from anesthetized cats.

3.2.2 Experimental methods

The experimental protocol of Sillito and colleagues will now be briefly described; it has been described in detail elsewhere (Sillito et al., 1993, 1994). Sillito et al. performed experiments on female cats anesthetized with halothane and paralysed with gallamine triethiodide. Visual stimuli were produced with a Picasso Image Generator (John Daugman) and displayed on a Tektronix 604 oscilloscope-like cathode ray tube. Extracellular recordings of single unit activity in the A laminae of the LGN were made with arrays of three electrodes, the tips of which were sufficiently separated from each other so the receptive fields of the LGN cells recorded from would be separated by between 1° and 4° (Sillito et al., 1994).

3.2.3 Visual stimulation

Receptive fields were initially hand-mapped and then more precisely located by reference to the timing of PSTHs generated by drifting precisely controlled bars, moving at various orientations, across the receptive field. Once the receptive field location and type (On/Off/X/Y) of all cells being recorded from had been determined, the cells were driven by computer-controlled visual stimuli. Two main types of visual stimuli will be discussed in this thesis: (1) full-field, drifting sine wave gratings; and (2) pairs of flashing or moving dots centered on the receptive fields of two of the cells. The analysis will focus on records obtained with the first type, full-field drifting gratings. Their arrangement and placing with respect to the receptive fields (RFs) of the cells being recorded from is illustrated in Figure 3.5.

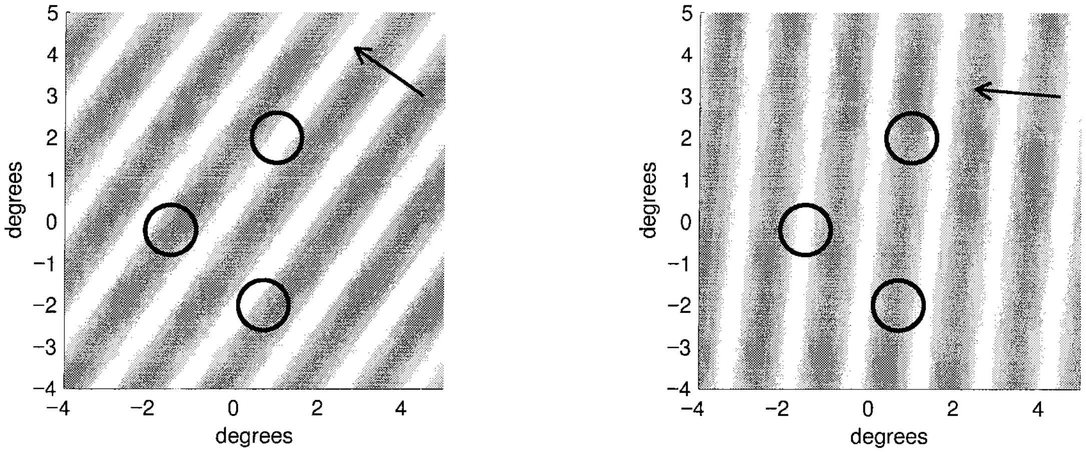

Figure 3.5: Diagram of the receptive fields (RFs) of 3 LGN neurons being recorded from. The same 3 receptive fields are shown in the two panels, but overlaid on two different full-field drifting sine wave gratings (used as stimuli). The direction of motion of the gratings is indicated by the arrows. Notice that as the direction of the grating changes, the relative phase of the stimulus on the different receptive fields also changes. In the left panel, the stimulus for the two rightmost RFs is out of phase, whereas in the right panel the same two RFs observe the stimulus in phase. The gratings used typically had a spatial frequency of 0.66 cycles per degree. Receptive fields are roughly 1° in diameter or less and are between 1° and 4° apart. (Sine wave gratings always move (perceptually) exactly orthogonally to their orientation; I will use "orientation" and "direction of motion" of the grating almost synonymously.)

Spatial frequencies of the gratings were in the range 0.66 to 1 cycle per degree; temporal frequencies were 2 or 3 Hertz. In a typical experiment, a set of 5 or 6 different orientations would be chosen, and a pseudorandomly interleaved sequence of presentations of these orientations would then be used to drive the cells. Each presentation of an oriented grating (henceforth called a *trial*) typically lasted 1.7 to 2.5 seconds, time enough for 5 modulations of the grating to drift over the receptive field of any one of the cells. An example of rasters obtained from one of the cells under these conditions is shown in Figure 3.6. The set of responses to many trials at one particular orientation, taken over a single experiment, will be henceforth called a *record*. The trials which composed the typically 5 or 6 records of an experiment – one record per orientation – were pseudorandomly interleaved with each other. A single set of records, the trials of which were interleaved with each other in the manner described, will be called an *experiment*. Experiments typically lasted on the order of 1 to 2 hours. Each pair of cells had data recorded from them over several experiments; different experiments used different sets of orientations and/or spatial or temporal frequencies. For some pairs of cells, additional experimental runs, using flashing or drifting spots or bars as visual stimuli, were also made.

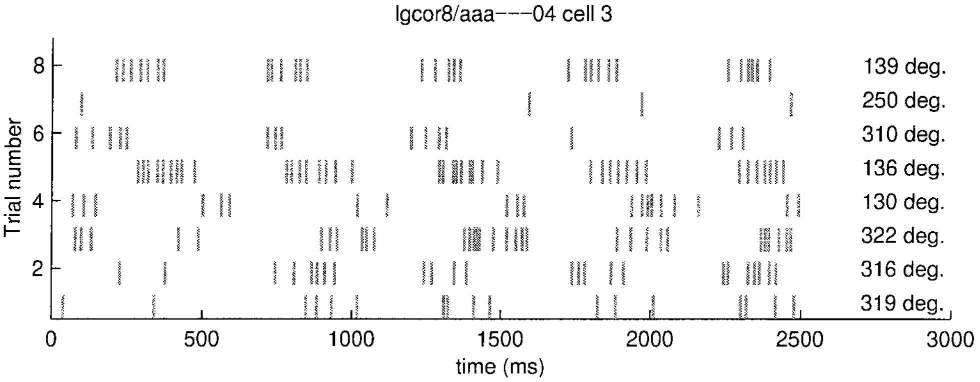

Figure 3.6: Responses to 8 successive trials in an experiment using a 2 Hz, 0.66 cycles/degree grating, drifting at various orientations. The orientation used in each trial is indicated at the right. Modulation of the responses at the temporal frequency of the grating, as described in Figure 3.1, are seen for all stimulus orientations. But as expected from Figure 3.5, different orientations result in response modulation at different phases. The grating orientation always remained constant throughout the length of each trial. The data was collected by Sillito et al. in anesthetized cats.

3.3 Correlated internal parameters

The cross-covariogram analyses done here and done by Sillito et al. (1994) are of periodically driven neurons whose spike trains have been broken into separate trials. As we shall see below, the correlations found in the data may have a component with a time scale *longer* than the time scale of the periodic driving and of the separation into trials. This section considers how these two characteristics (periodic driving and separation into trials) can affect the width of the peaks in a covariogram. We will explore the issue using two synthetic signals as an example.

Take two slowly varying, correlated signals $s_1(t)$ and $s_2(t)$. For convenience and without loss of generality, assume they both have mean zero. At any point in time $s_1(t)$ is approximately equal to $s_2(t)$, which is reflected by a peak at zero in the cross-correlation of the two signals, which we will call $c(t)$ and define as $c(t) = s_1 \otimes s_2$ (see Figure 3.7). The width of the peak is determined in part by how slowly s_1 and s_2 vary.

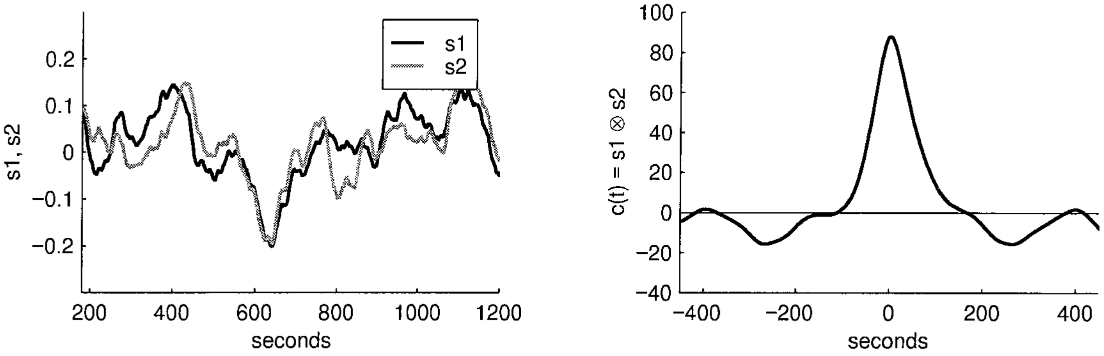

Figure 3.7: Two slowly covarying signals s_1 and s_2. On the right is the cross-correlation $c(t)$ of the two signals s_1 and s_2. Compare the width of the peak to the time scale of the variations on the left.

3.3.1 Splitting the data into trials

Let the two signals s_1 and s_2 be broken up into sections of equal length, and let us label these sections as different "trials". If the trials are shorter than the time scale of the correlations, and a covariogram of the different trials is built,[9] the width of the covariogram peak can no longer reflect the time scale of the correlations: the entire covariogram itself is only twice as wide as the trial length (see Figure 3.8). Though the peak in the covariogram is constrained to be no wider than the length of each trial, information about the longer time scale of the underlying correlations may still be obtained from calculating covariograms of $\{s_1^r\}$ and sets of "shifted" trials of signal number two: $\{s_2^{r+i}\}$. For example, if successive trials are separated by 21 seconds then the height of the peak in the covariogram of $\{s_1^r\}$ and $\{s_2^{r+1}\}$ will be proportional to the height of the correlogram $c(t)$ evaluated at $t = 21$ s. More generally (as long as trials are much shorter than the time scale of the correlations) the peak height of the covariogram of $\{s_1^r\}$ and $\{s_2^{r+i}\}$ will be proportional to the height of $c(iT)$, where T is the time in between successive trials. This allows using the trial

---

[9] The covariogram of two sets of signals $\{s_1^r\}$ and $\{s_2^r\}$ is $\langle s_1^r \otimes s_2^r\rangle - \langle s_1^r\rangle \otimes \langle s_2^r\rangle$. Each $s^r(t)$ is defined as zero outside the time of trial r.

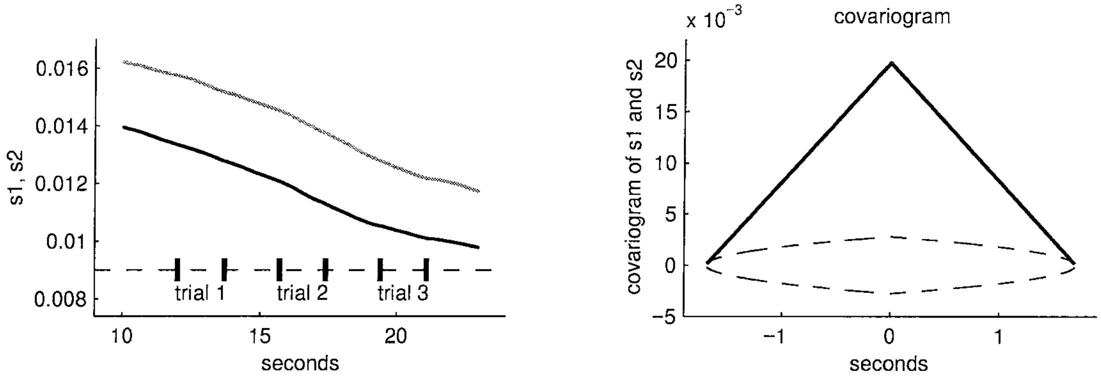

Figure 3.8: The left panel shows a short segment of the signals shown in Figure 3.7. The two signals will be broken into even shorter trials r, each 1.5 seconds in duration. (Trials may be separated from each other by a few seconds.) The right panel shows the cross-covariogram of the sets of trial segments of the data, $\{s_1^r\}$ and $\{s_2^r\}$. Dashed lines are significance limits, computed as in section 2.2: they are the 2σ limits of the expected variations if the two observed signals s_1^r and s_2^r were independent. The covariogram has a positive peak because the two signals are positively correlated. Each trial is 1.5 seconds long, so the covariogram is constrained to fall to zero at $t = \pm 1.5$, despite the fact that the timescale of the correlations of s_1 and s_2 is much longer. The shape of the peak is very close to triangular because the trials are short enough that s_1 and s_2 are almost constant within each trial; the covariogram thus has the shape of the cross-correlation of two boxcar functions, which is an isosceles triangle.

segment data to estimate the width of the peak in the original correlogram $c(t)$ (see Figure 3.9).

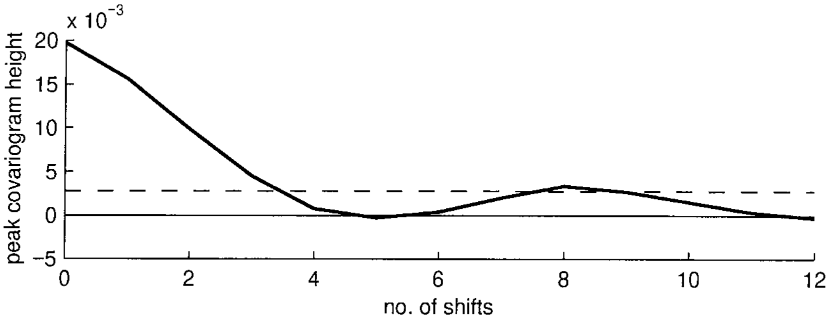

Figure 3.9: Peak height of the covariogram of $\{s_1^r\}$ and $\{s_2^{r+i}\}$ as a function of the shift i. The dashed line is the 2σ limit of being significantly different from zero. The number of shifts at which the peak height falls to within the significance limits (here approximately 4), multiplied by the time between successive trials (here 21 seconds), gives an estimate of the half-width of the cross-correlation peak of the underlying variables (approx. 84 seconds). This estimate should be compared with the width of the cross-correlogram peak in Figure 3.7.

3.3.2 Periodic modulation

The neuronal data that will be analyzed below was obtained from sensory neurons driven by a periodic stimulus. As described in Figure 3.1 and in section 3.2.1, the temporal profile of LGN cells in response to a drifting sine wave grating is that of a half-wave rectified signal. This is important, because the neurons were recorded extracellularly, and it is only when the neurons fire that information is obtained about their internal parameters. During periods of silence, the only information available is that they have not been driven to produce spikes.

Write the set of spike trains of cell 1 and cell 2 as $\{S_1^r(t)\}$ and $\{S_2^r(t)\}$, respectively, and the periodic stimulus for cell 1 as $p(\omega t + \phi_1)$ and for cell 2 as $p(\omega t + \phi_2)$ (both cells are driven by the same stimulus, but because of differences in their receptive field locations they may observe it at different phases– see Figure 3.5). The stimulus is periodic such that $p(x + 2\pi) = p(x)$. For times $t_o = (2n+1)\pi/\omega + \phi_1 - \phi_2$, where n is an integer, $p(\omega t + \phi_1)$ and $p(\omega(t+t_o)+\phi_2)$ have exactly opposite phases; due to the rectification of responses, this means that one of $S_1^r(t)$ and $S_1^r(t + t_o)$ must be zero, over all trials r and times t. Therefore the correlogram $C^r(t) = \sum_\tau S_1^r(t+\tau)S_2^r(\tau)$ evaluated at t_o will also be zero: $C^r(t_o) = 0$. This holds for all trials r, so the covariogram $V(t_o) = 0$ also. Thus $V(t_o) = V((2n + 1)\pi/\omega + \phi_1 - \phi_2) = 0$ for all integers n: that is, the covariogram is modulated with the periodicity of the stimulus. As an illustration of this, let

$m_1^r(t) = s_1^r(t) \cdot p(t)$ and $m_2^r(t) = s_2^r(t) \cdot p(t)$ represent multiplicatively modulated trials. The two sets of signals $s_1^r(t)$ and $s_2^r(t)$ are the same as were used in Figure 3.8. The modulating function $p(t)$ is shown in Figure 3.10 and represents the periodic driving stimulus. The covariogram of the modulated trials $\{m_1^r\}$ and $\{m_2^r\}$ is shown in Figure 3.11. The covariogram can be seen to be modulated at the stimulus driving frequency.

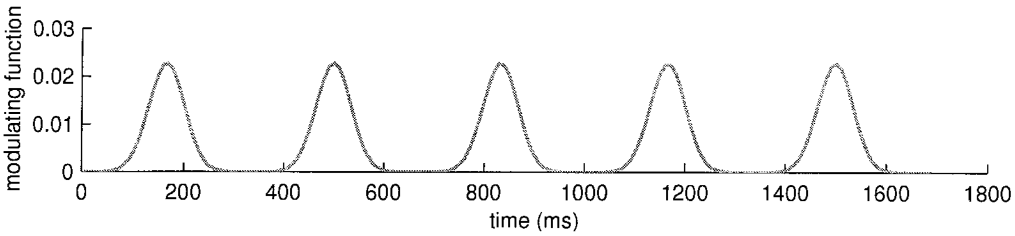

Figure 3.10: Periodic modulation function $p(t)$, frequency 3 Hz.

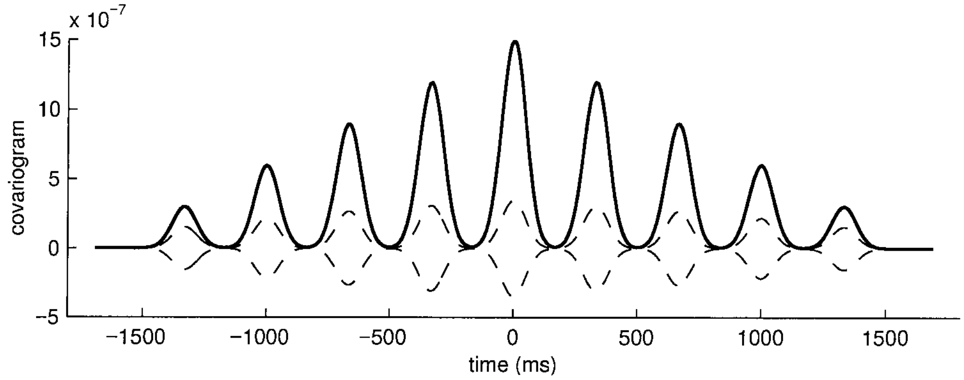

Figure 3.11: Covariogram (thick dark line) of the same trials used in Figure 3.8 but with the trials modulated by a 3 Hz periodic function. The resulting covariogram is modulated at the same frequency. The significance limits (shown as dashed lines) are also modulated. This is because the significance limits are computed based on the average signals $\langle m_1^r \rangle_r$ and $\langle m_2^r \rangle_r$, which have also been modulated by the periodic modulating function (see equation (2.3)).

In the neuronal data from the experiments the modulation is not a simple multiplicative one, and the spikes used to construct the covariograms are not simple linear functions of the underlying correlated parameters. The precise shape of modulations in the covariogram will not always be exactly like those seen in Figure 3.11. Nevertheless, the general features of the covariogram on Figure 3.11 are characteristic of the data, and the same method of shifting trials to determine the long time scale of the correlations that was described in section 3.3.1 will be used with the experimental data.

3.4 Experimental results

Twelve groups of recordings were made available for the analysis done here; this set of data was chosen by the experimenters (Sillito et al.) and thought to be a representative sample of their full data set.

Each of the 12 groups corresponds to a particular electrode *triad position*, at each of which data was taken during several different experiments. Of the 12 triad position recordings available for analysis, 5 had experiments in which there were significant covariations between two electrode sites. None of the experiments, out of all 12 triad positions, showed covariations between three sites simultaneously.[10] While the sample analysed here is only a partial one (5 out of a total of 19 pairs of significantly covarying cells recorded in the presence of cortical feedback (Sillito et al., 1994)), it is worth emphasizing that the sample was chosen by the experimenters, *before* the analysis done here, as one representative of the whole data set.

As described previously (section 3.2.3), by *experiment* I mean a single continuous data collection run, which typically lasted on the order of 1 to 2 hours. For the drifting gratings experiments, each experiment was composed of many different *trials*, each of which lasted on the order of 1 to 2 seconds. During each trial a drifting full-field grating at one of several possible orientations (typically 5 or 6 orientations) was presented to the animal. A similar amount of time elapsed between trials; during inter-trial times the visual stimulation screen was blank. The set of all responses to a single, particular orientation during an experiment will be referred to as a *record*. Thus the trials that compose the different records in an experiment were interleaved with each other, and all trials in a particular record used exactly the same visual stimulus.

Different experiments differed in the choice of orientations to be used, spatial or temporal frequency of the gratings, or in other parameters.[11]

---

[10] Simultaneous recording from three different sites is very difficult: often only two of the three sites had stable recordings in which single cells could be reliably identified.

[11] Many of the experiments analysed here used gratings that differed in orientations only by very small amounts (1 degree or less). It was thought that these small differences could be significant. However, the analysis of section 3.4.3 shows that all orientations produced covariograms with similar peaks. This conclusion is true both for experiments which used orientations within a small range and for experiments which used widely differing orientations.

3.4.1 Typical covariograms

Examples of typical covariograms obtained from from drifting sine wave grating records are shown in Figures 3.12 and 3.13. Note that each covariogram is constructed from a particular record. Since the stimulus is identical for all trials within a record, the covariations between the two spike trains obtained from two electrodes, evidenced by the covariogram peaks, cannot be attributed to variations in the stimulus.

A striking feature of these covariograms is that there is not only a central peak, near relative time $t = 0$, but there are also side peaks, spaced with the periodicity of the stimulus. In *all* cases where a significant covariogram peak was found near $t = 0$ axis, significant side peaks, spaced with the stimulus periodicity, were also found. These peaks cannot be explained by the mere fact that the stimulus is periodic. The example of Figure 3.14 shows a covariogram (made from synthetic spike trains) where the stimulus was periodic but only a central peak exists. The side peaks necessarily indicate that there is a component of the correlations with a time scale longer than the periodicity of the stimulus.

The observation of these side peaks is a major difference with the analysis reported in Sillito et al. (1994). When that analysis was done, it was thought that covariogram peaks necessarily represented spike timing correlations, and that these would not lead to peaks far from $t = 0$ (but see Abeles (1991)). Following this logic, the covariogram and JPSTH analysis was confined to a time window of $t \approx \pm 175$ ms. Most side peaks fall outside this window, and were therefore simply not observed. Figure 3.15 shows some typical covariograms as seen in a window confined to this small time interval.

However, the side peaks are very important. Their presence is intriguing, and clearly requires an explanation. It is also clear that unless these neurons are thought of as participating in synfire chains (Abeles, 1991), it is difficult to interpret the peaks far from $t = 0$ as spike timing correlation peaks. Most importantly, *whatever process explains the side peaks is also a good candidate process for explaining the central peak.* Such a parsimonious explanation, involving a single mechanism that accounts for both central and side peaks, will be put forward in this thesis.

As will be explained in this and the next chapter, the hypothesis that the resting membrane potentials of the two neurons being recorded from are covarying, and that they do so slowly (with a timescale of tens of seconds to minutes), can account for much of the data. A comparison of Figures 3.12 and 3.13 with Figure 3.11 already suggests that the timescale of the correlations

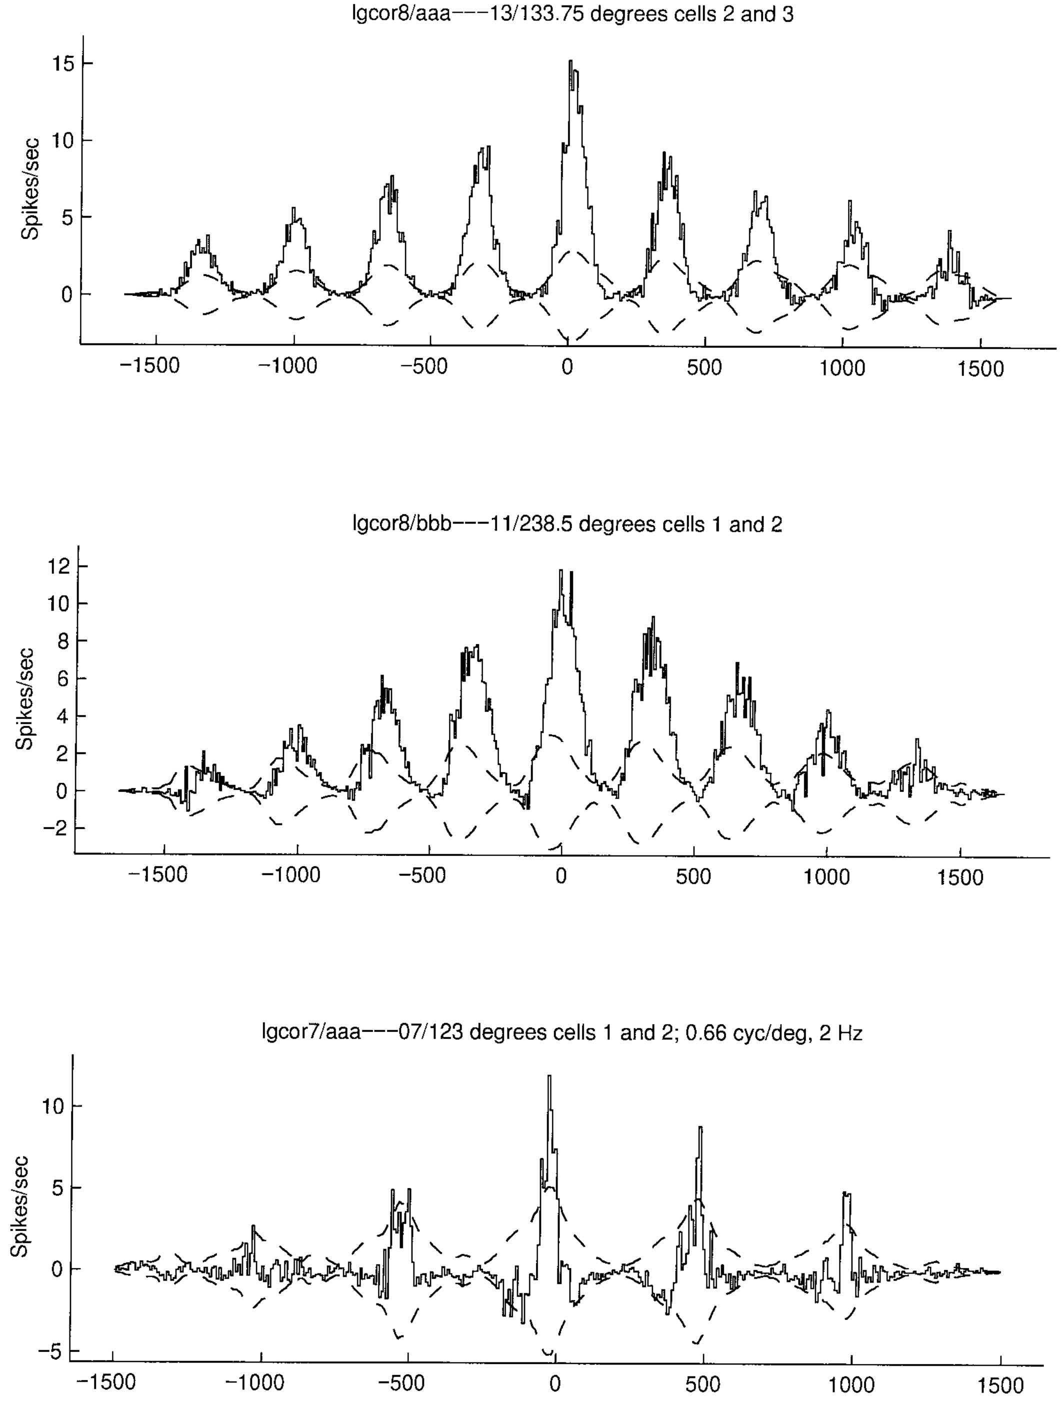

Figure 3.12: Covariograms of 3 different pairs of cells from 2 different animals. Time is in milliseconds. Dashed lines are significance limits, computed as in equation (2.3). The striking presence of "side" peaks off $t = 0$ indicate that the correlations have a strong component with a time scale longer than the stimulus period (which in all cases was either 333 or 500 ms). Section 3.4.2 will show that the correlations actually have a time scale on the order of tens to hundreds of seconds (not milliseconds). The original report on this data (Sillito et al., 1994) was limited only to the $t = \pm 175$ ms window of these covariograms, and thus did not observe the side peaks. The presence of side peaks suggests an interpretation of the covariograms significantly different to that in the original report. **Normalization of covariograms:** In this chapter, covariograms, which have natural y-axis units of spikes^2 per bin, were normalized by dividing by the bin width and by $\sqrt{\langle n_1^r \rangle \langle n_2^r \rangle}$, where n_1^r and n_2^r are the total number of spikes fired by cell 1 and cell 2 in trial r, respectively, and the average is taken over trials. This normalization gives final y-axis units of spikes per second and allows comparing covariograms taken from experiments with different numbers of total spikes fired per trial. Note that the normalization only changes the y-axis labelling, it does not change the shape

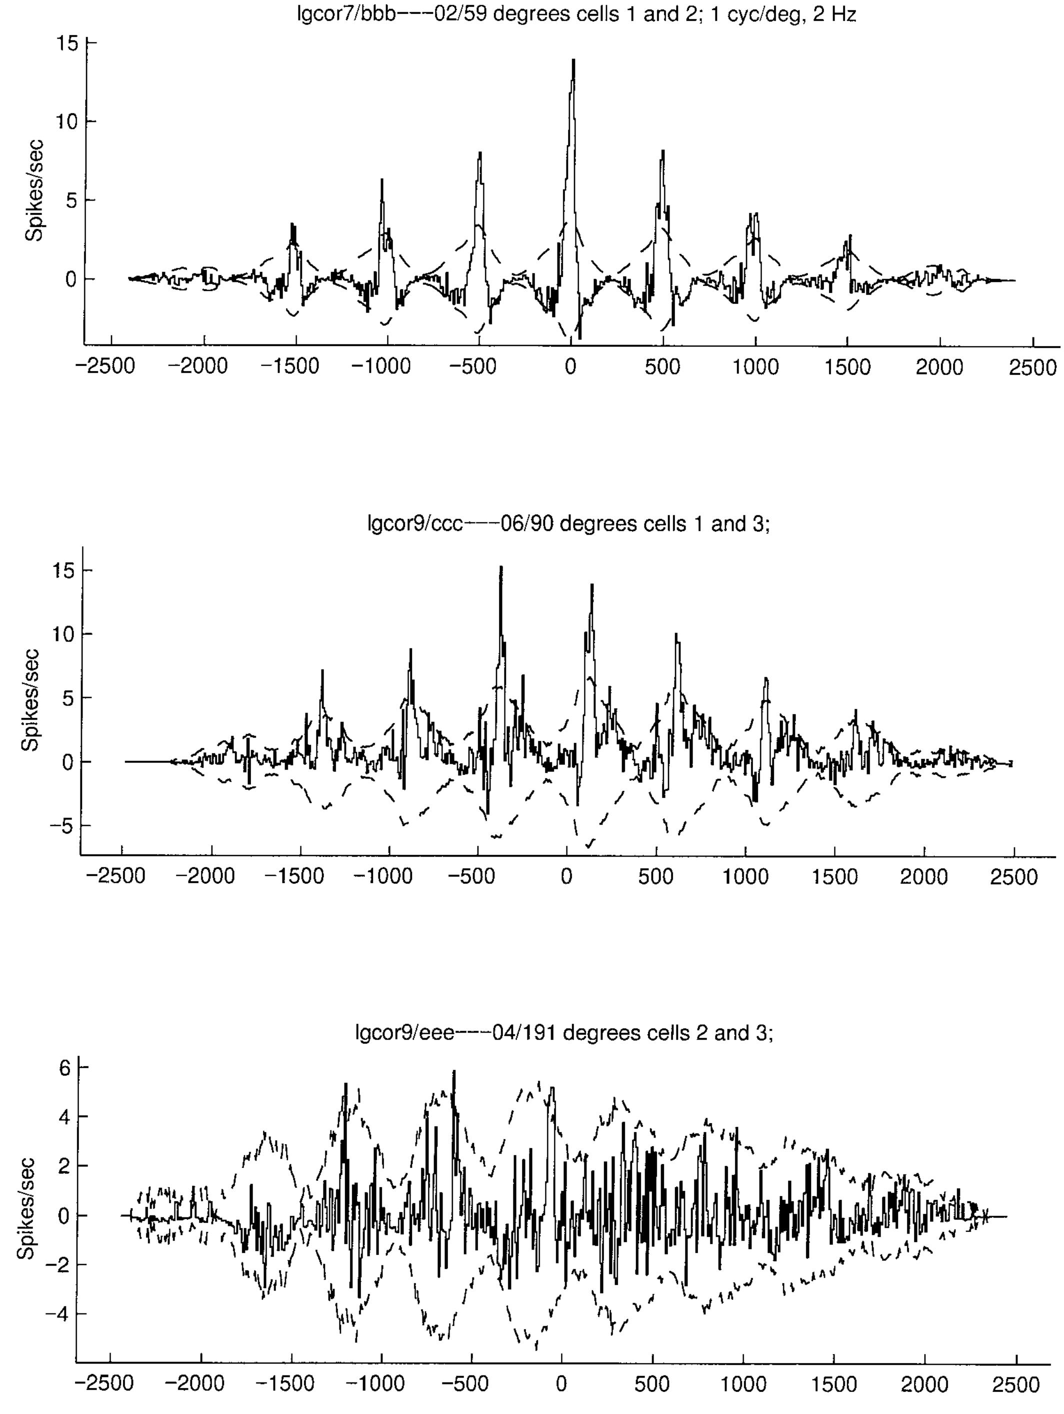

Figure 3.13: Covariograms of 3 different pairs of cells from one of the same two animals as Figure 3.12 and from a third, different animal. Time is in milliseconds. The bottom covariogram is from the same animal as the middle covariogram, and serves here as an example of two cells that were not significantly correlated. Dashed lines are significance limits, computed as in equation (2.3).

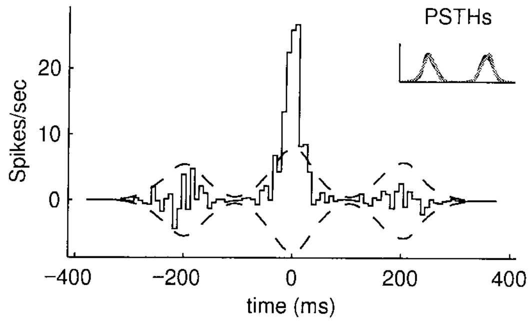

Figure 3.14: Even when the stimulus is periodic, a central peak does not imply the presence of side peaks: this Figure, which should be compared with Figures 3.12 and 3.13, is an example of a central peak without side peaks. Though the central peak is clearly significant, no significant side features are seen. (Dashed lines are significance limits). The spike trains used here are synthetic, created with the techniques of section 2.7.3. Inset are the PSTHs of the two simulated cells, showing the periodic nature of their mean responses; axes in the inset are 400 ms wide and 90 Hz tall.

is longer than the length of a single trial. The following two subsections will characterise the covariograms in terms of their timescale and orientation-dependence. Section 3.5 will show that the covariograms can be well reproduced in terms of latency and excitability correlations. And chapter 4 will show that they can be well modeled in terms of (slow) covariations in the resting membrane potentials of the two neurons.

3.4.2 What is the time scale of the correlations?

The "shifting" method described in section 3.3.1 will be used to determine whether the correlations have a long time scale. Figure 3.16 shows how shifting the responses of one of the cells by one or more trials affects one of the covariograms. That is, the Figure shows covariograms made using the spike train sets $\{S_1^r(t)\}$ and $\{S_2^{r+i}(t)\}$, where r is the trial number within a particular record of trials at a single orientation, and i is one of 0, 1, or 2. Though successive trials at one particular orientation were not always separated by a fixed amount of time (since different orientation trials were pseudorandomly interleaved with each other), the trials are nevertheless in an orderly sequence, and the average time difference between trials r and $r + i$, where i is the number of trial shifts, is indicated above each covariogram. It is clear that the cells remain correlated on a time scale much longer than the length of individual trials (which is 1.6 seconds).

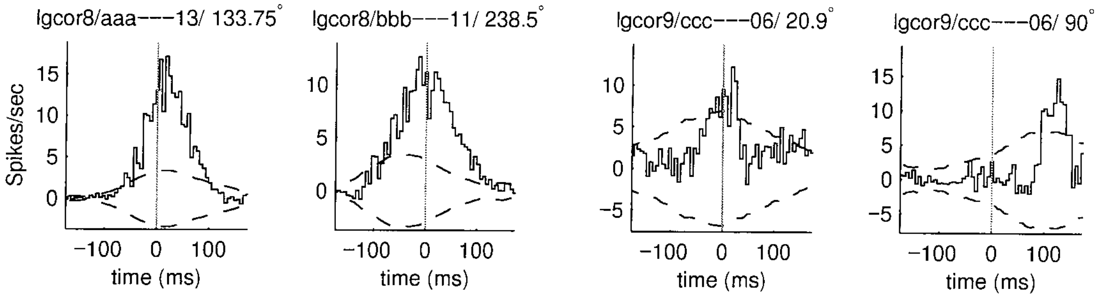

Figure 3.15: Typical covariograms, seen *only* in the window $t = \pm 175$ ms, as in the original report (Sillito et al., 1994). Most side peaks lie outside this window, and were simply not observed. This Figure should be compared to Figures 3.12 and 3.13, where the full covariograms are shown. (Note that some peaks wholly off $t = 0$ can nevertheless still be seen within the limited window. Vertical grey lines indicate the $t = 0$ axis.) The two rightmost covariograms are from the same experiment, and thus are of trials interleaved with each other; one used a grating at 20.9°, the other a grating at 90°. The time difference between the centers of these two peaks is exactly predicted by the difference in the relative phase of the PSTHs of the two cells.

It is important to note that trials were separated by periods of "silence" (blank stimulus) lasting approximately two seconds: whatever mechanism maintained the cells correlated over several trials cannot depend on brisk levels of activity in the visual pathway, since during the blank stimuli periods most of the the visual pathway is inactive.

Figure 3.17 shows, using the same data as in Figure 3.16, how the peak height of the cross-covariogram changes as a function of the number of shifts, i. (In the Figure the x-axis is labelled with the corresponding average time difference between trials, in seconds.) Plotted on the same Figure are the fall-offs in peak height as a function of shift for the *auto*-covariograms of the two cells involved. The fall-offs for the auto-covariograms follow the same time course as the fall-off of the cross-covariogram. On log-linear axes the peak height fall-offs are well fit by a straight line: the thick grey line is the least squares fit to all three data lines (autocovariogram of cell 2, autocovariogram of cell 3, cross-covariogram of cells 2 and 3). Based on this fit, the estimated time for the peak height to fall to within the significance limits is 42 seconds; the exponential decay time constant of the fit is 21 seconds.

Figure 3.18 summarizes the result of such fits, applied to all the data records where there were significant correlations. The left panel shows a histogram of the estimated times to fall below

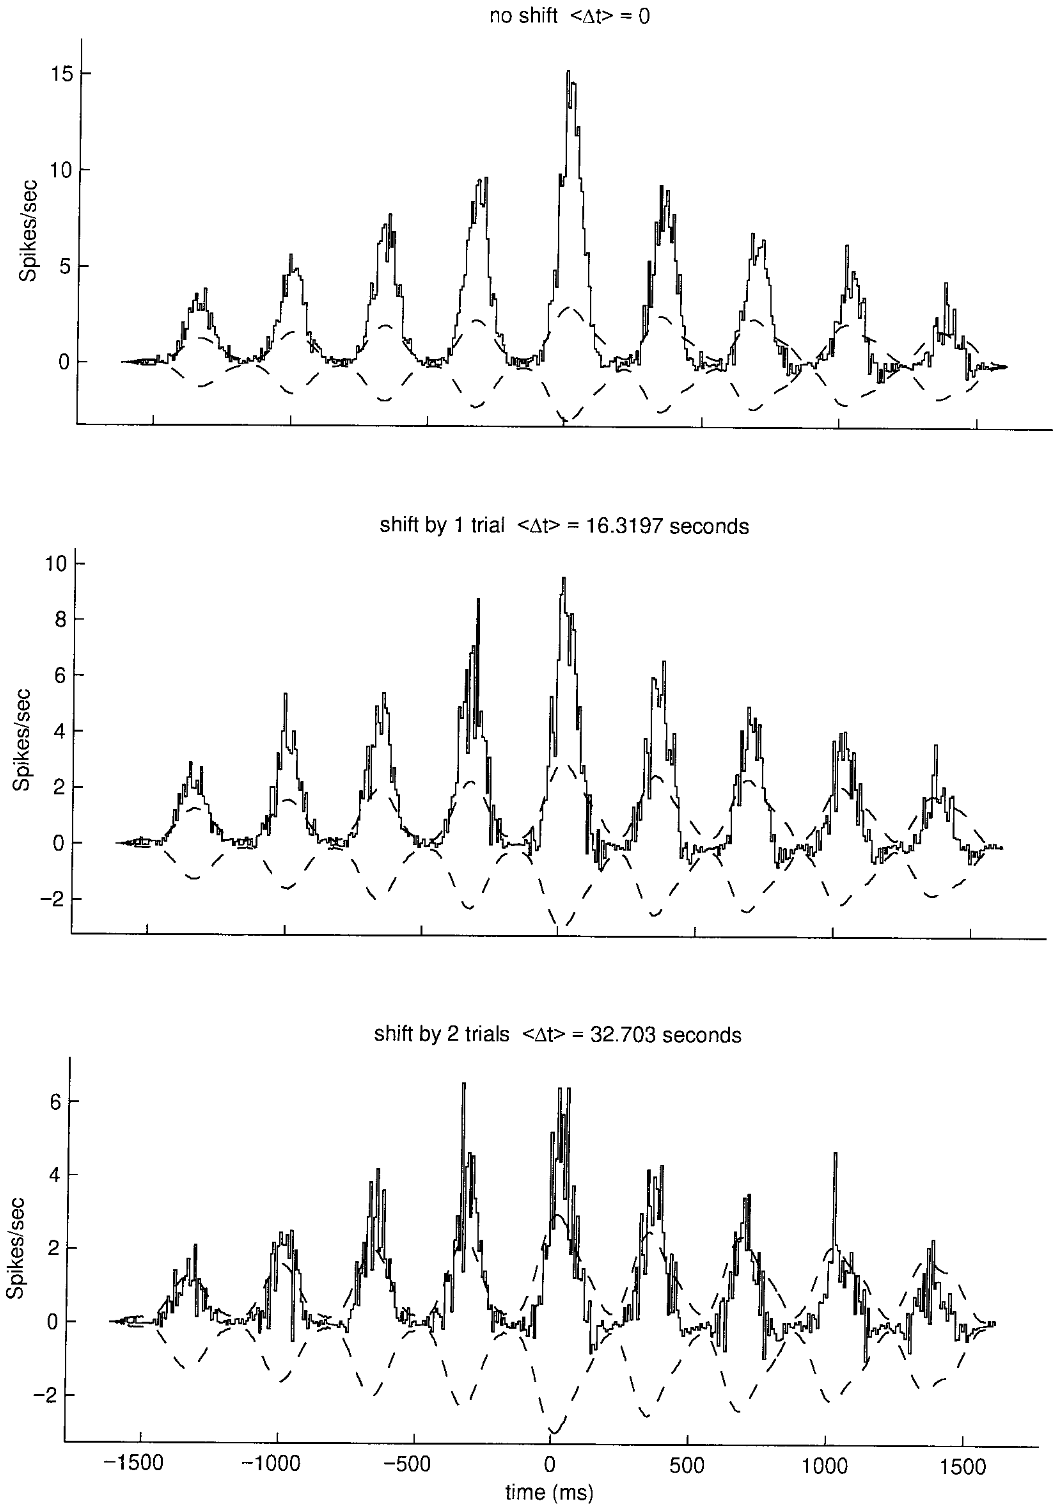

Figure 3.16: An example of an unshifted cross-covariogram and the results of shifting by 1 and 2 trials, respectively. Peak heights get smaller (note the different scalings on the y-axes), but remain significant for times on the order of tens of seconds. (Data record `lgcor8/aaa---13/133.75`°, cells 2 and 3.)

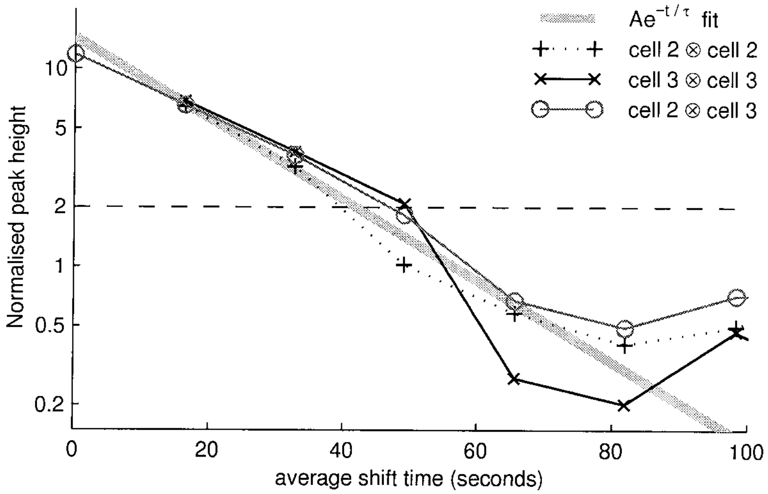

Figure 3.17: Fall-off in peak height as a function of shift for the cross-covariogram and autocovariograms of the same data record as in Figure 3.16. Note that the **y-axis** is on a **logarithmic scale**. Peak heights have been normalised by the maximum value of the standard deviation (σ) of the null hypothesis. (All covariograms in this thesis display 2σ as dashed lines to represent significance limits– see, for example, the dashed lines in Figure 3.16. Thus, a value of 2 on the graph above indicates the point where the maximum peak height is equal to the maximum significance limit.)

significance limits; the right panel shows a histogram of the exponential fit time constants. (The time to fall below significance limits depends not only on the exponential time constant but also on how strong the correlations at shift $i = 0$ are, so the two times are not proportional to each other.) The mean time to fall to significance limits is 41 seconds; the mean decay time constant is 64 seconds.

In conclusion, a strong component of the correlations has a time scale on the order of tens of seconds.

3.4.3 Does the strength of the correlations depend on the orientation of the stimulus?

As described in section 3.2.3, each experiment consisted of a set of trials using pseudo-randomly interleaved gratings of different orientations. Figure 3.19 shows covariograms from two different experiments. Each column corresponds to an experiment; the two rows are covariograms of two different orientation records from the experiment. As can be seen, both records from each experiment produce similar peaks. (The peaks in the two records are slightly shifted in time with respect

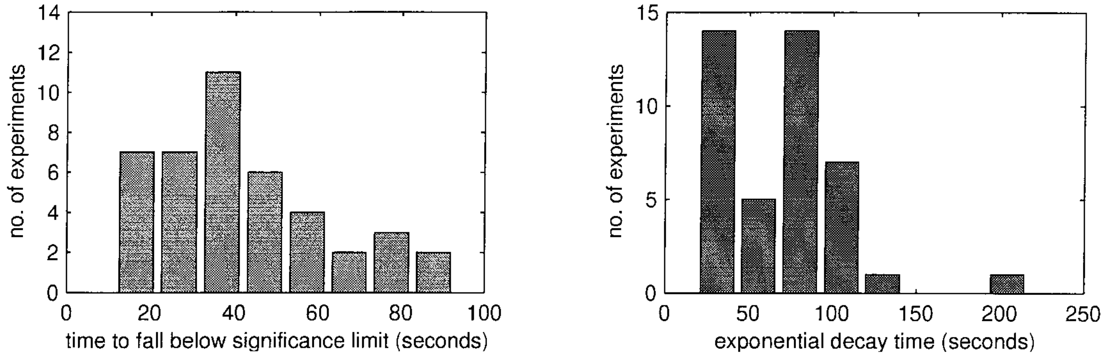

Figure 3.18: Histograms, across all experiments where significant correlations were found, of the result of shifting the trials and fitting the fall-offs in normalised peak heights. The left panel is a histogram of the estimated times to fall below significance limits. The mean time to fall below the significance limits is 41 seconds; its standard deviation is 16 seconds; its smallest value is 17 seconds. The right panel is a histogram of the estimated exponential decay time constants. The mean of the decay time constants is 64 seconds; their standard deviation is 46 seconds; and their smallest value is 19 seconds.

to each other. The shift follows the shift in the relative phases of the two cells' PSTHs.) The two orientations in the right column are 180 degrees apart. That is, they represent the same orientation but opposite directions of drift. Both directions of drift produce similar peaks.

These covariograms suggest that the height and strength of the peaks are independent of the orientation of the stimulus. The stimulus orientation's main effect is to determine the position of the peaks. An experiment suggesting the opposite conclusion is shown in Figure 3.20. Again, these are covariograms of data collected during the same experiment. The covariogram for data taken at 21.25° grating orientation has a much stronger peak than the one taken at 20.25°; but the magnitude of this difference must be quantified. We must ask whether the difference between the two covariograms is large or small compared to the expected measurement error. The fact that the orientation difference between the two records with in Figure 3.20 is only 1 degree, a difference well within the tuning width of cortical V1 cells and within human, monkey, and cat psychophysical limits (Lee et al., 1997; Sprague et al., 1993; Vogels and Orban, 1990), prompts careful examination of this experiment.

The null hypothesis that we will consider here is that all orientations produce the same peak covariogram heights, and that differences between them in peak heights are simply due to measurement noise. We must estimate the distribution of peak heights that such a null hypothesis

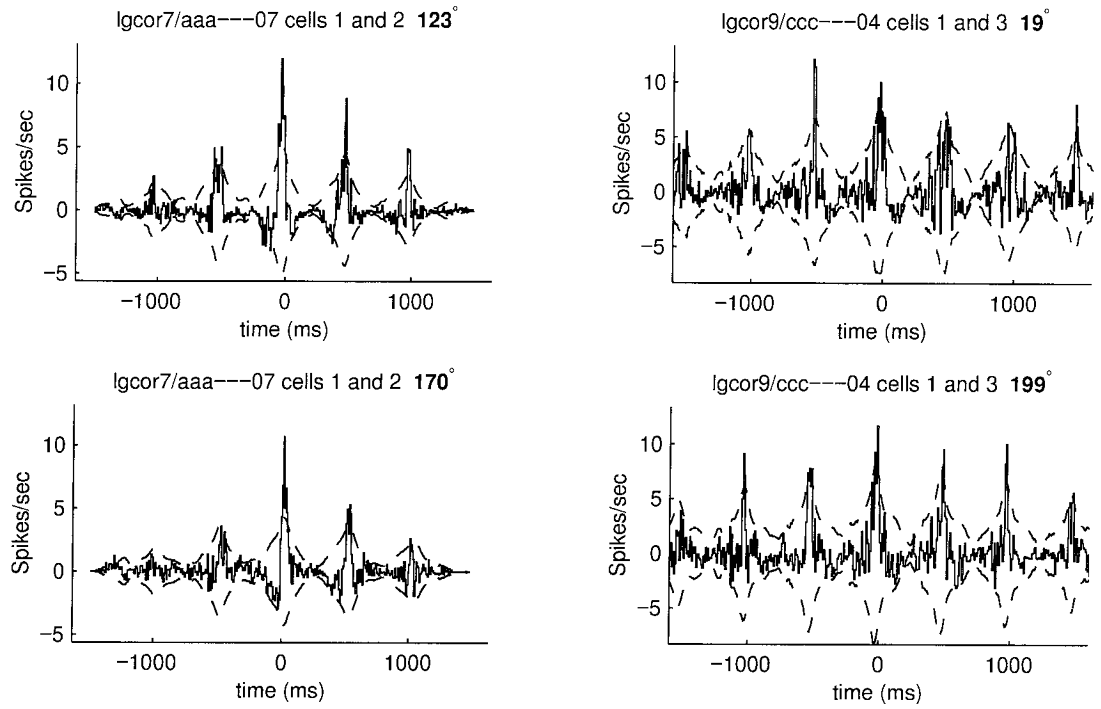

Figure 3.19: Covariograms at various orientations. Each column shows the covariogram of two records from the same experiment. (As explained in the text, trials for the records for the two covariograms in each column were pseudo-randomly interleaved with each other.) The left column shows covariograms from two different orientations; the right column shows covariograms from two opposite directions (the difference in orientation is 180°). The peaks in the two covariograms of each column are similar to each other, despite the differences in orientation and direction.

would produce, and evaluate how probable or improbable it is to find, in an experiment, a peak height the size of that shown in the middle and right panels of Figure 3.20. Estimating the distribution of peak heights will be done using a variation of the bootstrap technique (Press et al., 1992). To do this, all trials from all orientation records in an experiment are shuffled,[12] and then divided into as many sets as there were orientations in the original experiment. In the experiment shown in Figure 3.20, the covariogram of trials at 21.75° was chosen as the one with the biggest peak *after* covariogram analysis of all orientations in the experiment. Following precisely the same strategy, after the shuffling we examine the covariograms of all the sets and choose the one with the largest peak: this peak height then represents a sample data point, that of a maximum peak height as obtained under the null hypothesis. (Notice that the trials for each of the shuffled sets

---

[12] Note that unlike the "shifts" described in section 3.4.2, the spike trains of the two cells recorded in any trial are always kept together: it is entire trials, involving the responses of both neurons, that are used as the units to be shuffled here.

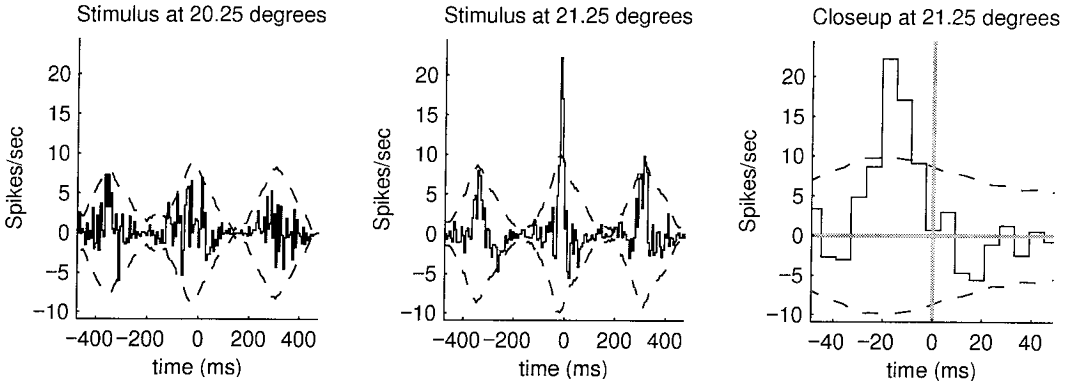

Figure 3.20: Two covariograms from the same experiment but of records at slightly different orientations. Left: stimulus at 20.25°. Middle: stimulus at 21.25°. The peak at 21.25° is much greater than any hints of a peak at 20.25°, suggesting that these two records lead to a statistically significant result. However, this difference must be compared to the expected measurement error. The significance limits, shown as dashed lines, give a rough estimate of the magnitude of the expected error. In the figure the significance limits are shown around $y = 0$; other mean values would displace the significance limits vertically. Right: closeup of the middle panel. Interestingly, the closeup shows that the peak is not centered at exactly $t = 0$ but is in fact ≈ 18 ms away from that. The two neurons being recorded from cannot then be strictly thought of as synchronized. The 18 ms time difference roughly corresponds to the phase difference in the PSTHs of the two cells. Experiment `lgcor9/ccc---08`.

Figure 3.21: The two stimuli which produced putatively significantly different covariograms. To this author's eyes, these stimuli are identical.

were not taken from any particular orientation.) Many such shuffles and covariogram peak height evaluations then generate a histogram estimating the distribution of maximum peak heights under the null hypothesis.[13]

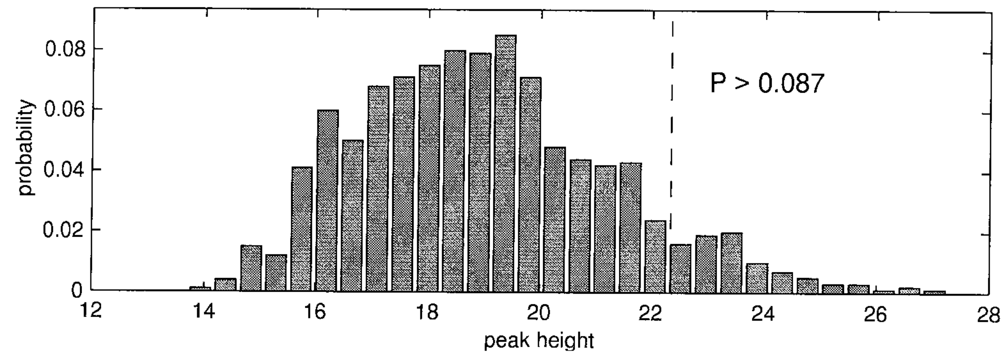

Figure 3.22: Distribution of the maximum peak height out of six covariograms, as obtained by the shuffling procedure described in the text, under the null hypothesis that all orientations produce peaks of the same height. The dashed line indicates the peak height of the record #lgcor9/ccc—08 at 21.75°.

Figure 3.22 shows the distribution of maximum peak heights obtained after 1000 random shufflings. The dashed line indicates the maximum peak height of the original record at 21.75°. In more than 8.5 per cent of the random shufflings, one of the six covariograms had a peak of that height or greater. The probability of finding a peak of that height or greater simply by chance is high.

For each experiment, let P be the probability (under the null hypothesis of all orientations producing peaks of the same height) of finding a maximum peak height greater than the largest peak in the original experiment. For the experiment shown in Figures 3.20 and 3.22, $P = 0.087$. The next Figure, 3.23, shows the values of P obtained over all experiments where significant covariations between pairs of cells existed. If the null hypothesis is always correct, the distribution of P should be a uniform one– while if the null hypothesis is always incorrect, the distribution should be markedly skewed, with many small values, since a small value would indicate that one orientation in the experiment had a peak height unexpectedly greater than any of the others.

---

[13] One must ensure that no latency correlations (due to including, within a single set, responses from many different orientation trials) have been introduced into the shuffled data. This was done by first shifting all responses in time so that the PSTHs of both cells, at all orientations, lined up. Note that these time shifts of the data do not change the peak heights of the original covariograms at all: the covariograms are merely displaced on the t-axis. After this time shift, trials from different orientations can be grouped into a single set without introducing any latency correlations, and the shuffling procedure proceeded as described above.

But the smallest value found was 0.059, (not small enough to be declared significant), and the distribution of P is fairly close to uniform.

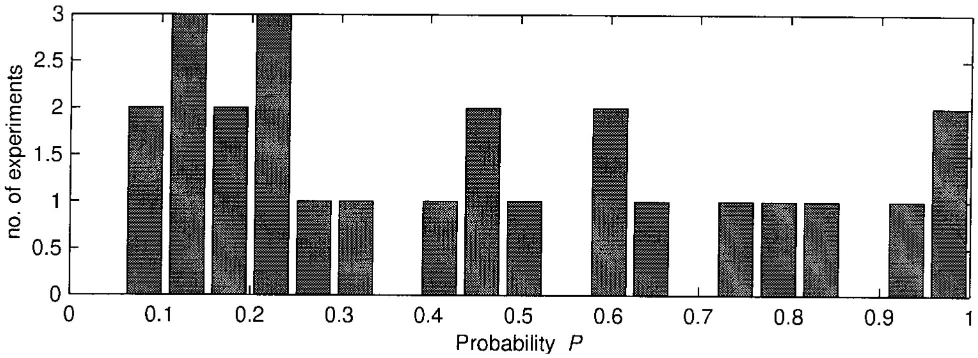

Figure 3.23: For each experiment, P is the probability of finding a maximum peak height bigger than the largest one found in the experiment. P is evaluated under the null hypothesis that all orientations generate similar covariogram peak heights. The graph above shows a histogram of the P's found in the experiments that had significant correlations between pairs of cells. No experiment had a P value of less than 0.05, and only two experiments had P values of less than 0.1. If the null hypothesis were always correct, all bins in the histogram above would be expected to have roughly the same height.

In conclusion, the orientation of the drifting grating used as a stimulus in these experiments does not significantly affect the height of the peaks in the covariogram.

3.4.4 Are responses to one orientation correlated with those at another?

Correlation strength is not influenced by the orientation of the stimulus (section 3.4.3), and the time scale of the correlations is on the order of tens of seconds (section 3.4.2). During these experiments, several different orientation trials were typically presented as stimuli, 'interleaved with each other, within any given 10 seconds. This leads to asking whether responses to the different orientations were correlated with each other. Figure 3.24 shows some example covariograms that demonstrate that responses at one orientation indeed often covaried with responses at a different orientation.

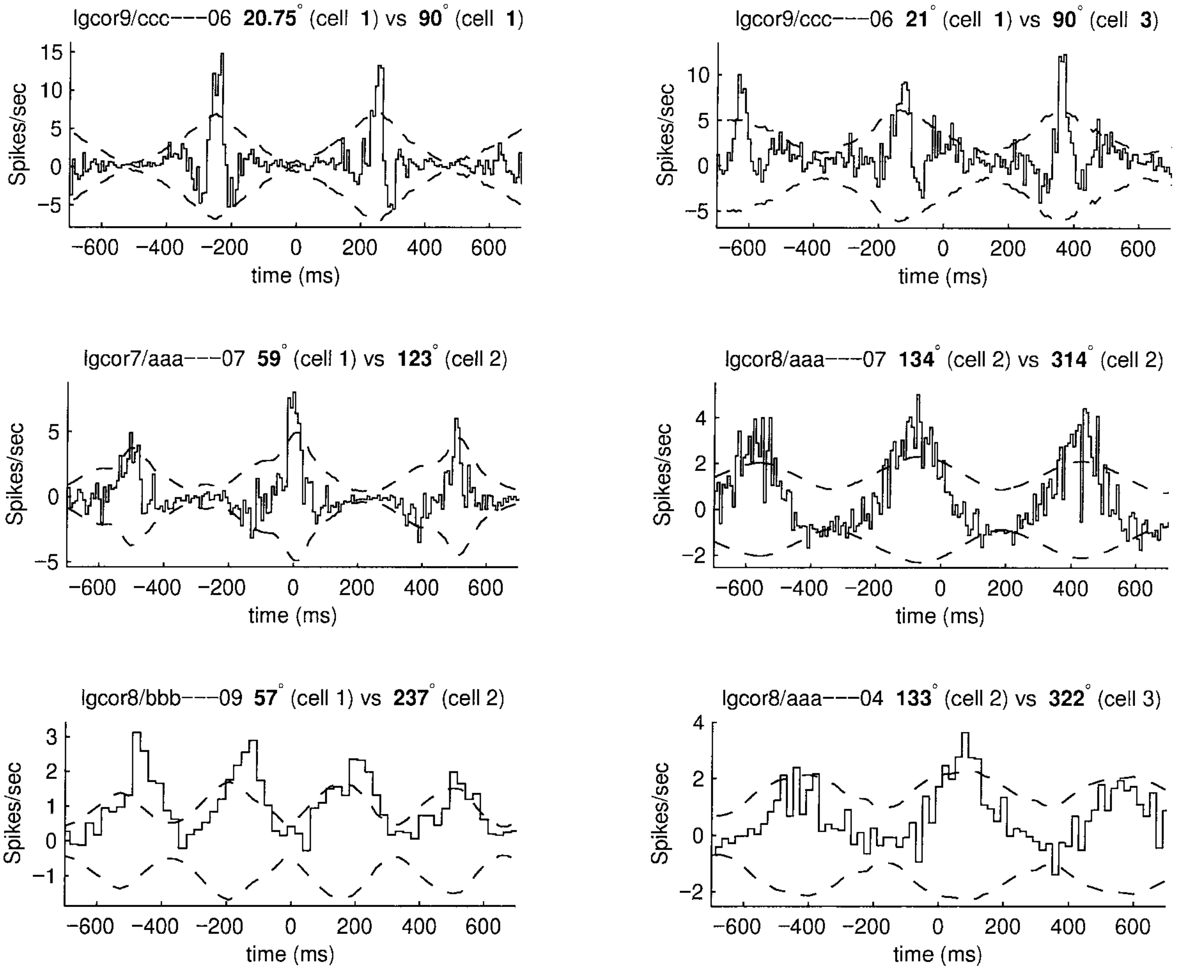

Figure 3.24: Covariograms of responses to one orientation vs. responses at another. The top left and the middle right covariograms use data from a single cell; all others use data from two cells. The bottom covariograms were made using large bins for greater clarity.

3.5 Latency and Excitability analysis

The analysis described in the previous sections dealt with the time scale and stimulus selectivity of the observed correlations. This section asks what kinds of correlations are present in the data, in terms of the types of correlations described in chapter 2. This will be answered by using the methods of section 2.4 (computing excitability estimates and performing latency searches).

Motivation for doing this type of analysis came from studying the raster plots of the responses. Figure 3.25 shows the raster plot for the experimental record whose covariogram is shown in the top panel of Figure 3.13. Data from the same experiment was used for Figure 2a in Sillito et al. (1994). In Figure 3.25, it can be seen that the cells go through relatively long periods of high firing rates (e.g. a period of about 50 sec. near start time 200 secs), alternating with periods of low firing rates (e.g. trials near start time 315 secs). Both cells go together through the same periods. The latency of the responses is also seen to covary over trials. As explained in chapter 2, such covariations, called latency and excitability correlations, lead to peaks in the covariogram of the two cells. Quantifying the size and shape of the peaks they contribute is the purpose of the excitability estimates and the latency searches.[14]

Under the experimental conditions used in these experiments, background firing rates were very low. Thus, for the excitability esimate the background firing rate gain parameter β^r (see section 2.4.1) was always set to zero and only one gain parameter, ρ^r , was estimated for each cell per raster r. This parameter was set to be directly proportional to the total number of spikes fired by the cell in that raster. In addition, one latency shift parameter for both cells, Δt^r, was estimated based on the latency search method (section 2.4.2). These sets of parameters, $\{\rho_1^r\}$, $\{\rho_2^r\}$, and $\{\Delta t^r\}$, were then used to predict the covariogram of two model cells that have PSTHs matching those of the real cells and which covary only in their latencies and excitabilities (but not in their individual spike times).

Figure 3.26 shows the result of such predictions, together with the original covariograms, for data from three different experiments. The predictions are remarkably similar to the experimental

---

[14] In their paper introducing the JPSTH technique, Aertsen et al. (1989) stated, referring to slow changes in mean firing rate such as those seen in Figure 3.25, that "careful screening of the spike train data for such nonstationarities is mandatory." They continued to say that "if necessary (and possible), one may use the procedure of "slicing" the spike trains into different sections, each one periodically stationary, and compare the results across sections." "Slicing" the data into subsets of trials with roughly similar mean firing rates has the objective of eliminating the excitability correlation contribution to a covariogram, and will be done, for completeness, below. However, the methods described in chapter 2 are proposed as more quantitative and complete procedures to address the same issue. In the original report by Sillito et al. (1994), the issue was not addressed.

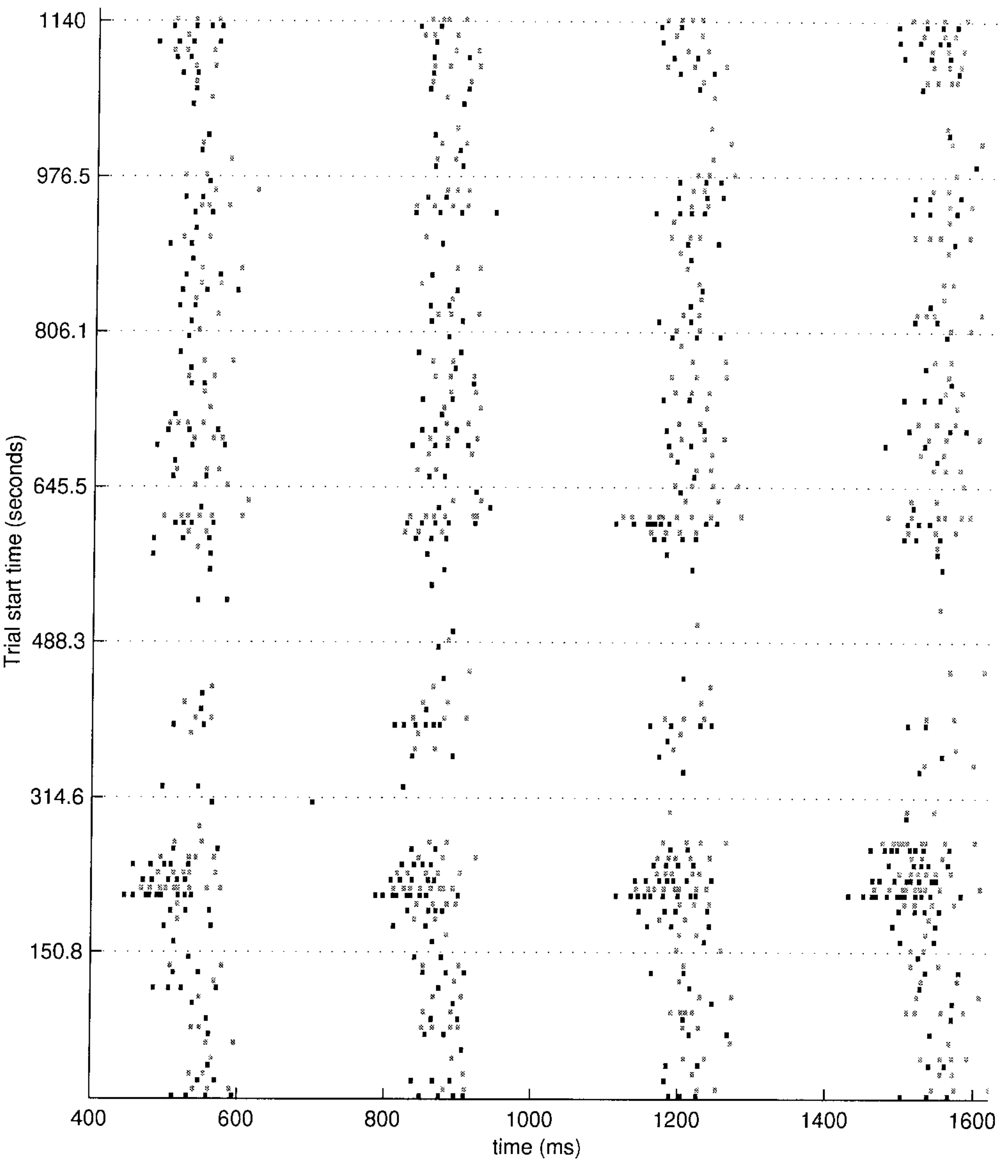

Figure 3.25: Rasters for the top panel of Figures 3.13 and 3.26, record #lgcor8/aaa—13 at 133.75°. It is easily seen that the cells go through long periods of high excitability (many spikes), alternating with long periods of low excitability (few spikes– e.g., compare the large numbers of spikes fired in response to trials near start time ≈ 200 seconds, compared to the very few spikes fired in response to trials near start time ≈ 315 seconds.). Since the excitabilities of both cells covary, this results in excitability correlation peaks in their covariogram. The latency of the response covaries also. One cell is shown in black, the other in grey; trial start time is shown on the y-axis.

covariograms, which shows that in order to account for the covariogram peaks, only 2 global parameters (ρ^r and Δt^r) need to covary between the two cells, and strongly suggests (though does not prove) that the correlations seen in the original data are not spike timing correlations, since the covariograms can be very well reproduced *without* spike timing correlations.

The predictions are based partly on direct estimates from the data (the excitability gain parameters) and partly on fits to the original covariograms (the latency search). Why are the predictions *so* good? It is possible that even fewer than two parameters per raster are covarying in the cells: the quality of the predictions suggests that two parameters per raster may have been enough to overfit the data. Another reason why the fits are so good is that the original covariogram peak shapes are very well matched to the shapes of the PSTHs of the two cells. As explained in chapter 2, the shape of peaks caused by latency and excitability correlations is largely determined by the shape of the PSTHs (in contrast to the shape of peaks due to spike timing correlations, which can be much more arbitrary). The shapes being so well matched again suggests that the correlations are overwhelmingly of the latency and excitability type.

So far, only covariograms have been shown, but the results also hold for JPSTHs. The same latencies, excitabilities, and PSTHs that were used to make the prediction shown in in the top panel of Figure 3.26, were used, in conjunction with the methods described in section 2.7, to generate a synthetic pair of sets of spike trains (one set for each of the two cells involved in the covariogram). The JPSTH of the synthetic sets, and the JPSTH of the experimental spike trains, are compared in Figures 3.27 and 3.28. Their similarity indicates that the main features of the experimental JPSTH have been captured by the latency and excitability description.

The excitability estimate and latency search are intended to be used for separating out the various components of covariograms. With respect to the experimental data, the result of those anlyses so far have argued the the latency and excitability correlation contributions are the major components of the covariograms. If this is true, direct measurements of excitability and estimated latencies should show characteristics similar to those of the peaks in the experimental covariograms. In particular, their time scale should be similar. Figure 3.29 shows the excitabilities of two neurons over the course of an experiment.[15] These are the same neurons whose covariogram was shown in the top panel of Figure 3.26. The time scale of the variations in excitability is similar to that of the correlations: it is on the order of tens of seconds (section 3.4.2). A similar plot show-

---

[15] In this case, since the background firing rate was almost zero, "excitability" is synonymous with "mean firing rate".

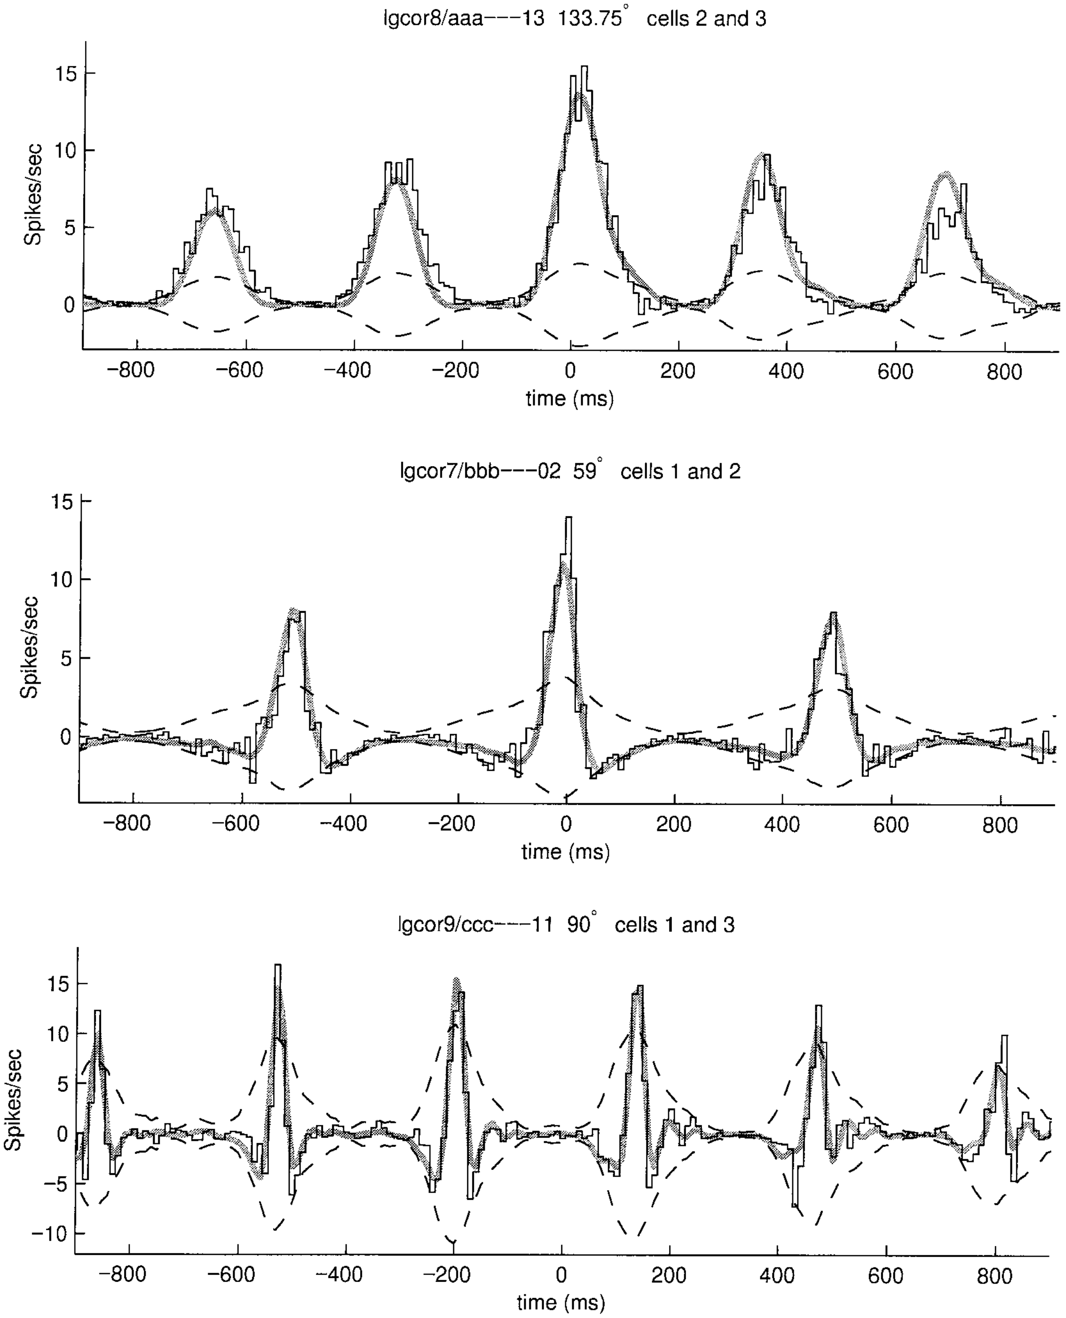

Figure 3.26: Covariograms of experimental data and predictions based on the latency and excitability analysis. The covariograms are shown as the thin dark lines (binned); the predictions are the thick grey lines. Dashed lines indicate significance limits. The top covariogram uses data from cells in the *tonic* mode (neither cell fired bursts during the experiment). The middle covariogram uses data from cells in a *mixed* mode (the cells sometimes fired in bursts, sometimes tonically); the bottom covariogram uses data from cells in the *burst* mode (both cells fired almost always in bursts). The fits to the data, shown as the grey lines, are made using models where the only covariations between the two cells are in the latencies and excitabilities of the responses, not in individual spike times. The peak shapes predicted by such models are closely related to the PSTHs of the two cells' responses (in contrast to 'the peak shapes allowed by spike timing interactions, which can be much more arbitrary). As can be seen, the fits are very good.

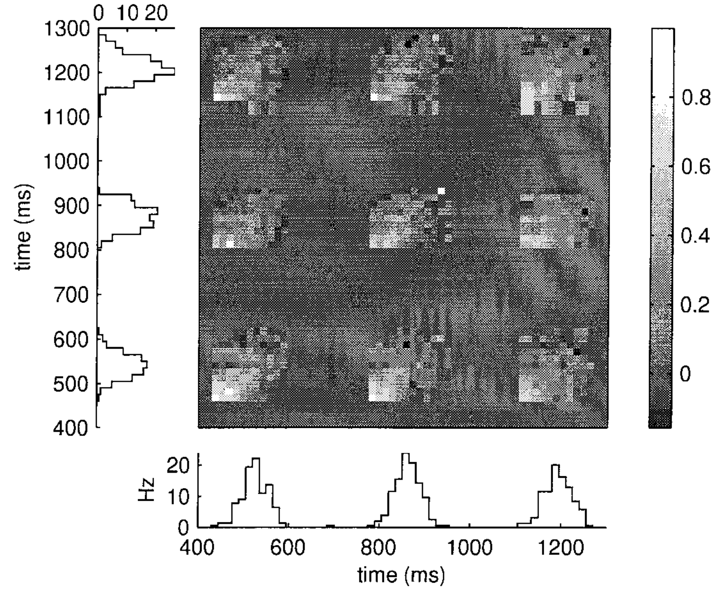

Figure 3.27: JPSTH of the experimental spike train data used in the top covariogram of Figure 3.26. (Data record `lgcor8/aaa---13 133.75° cells 2 and 3`).

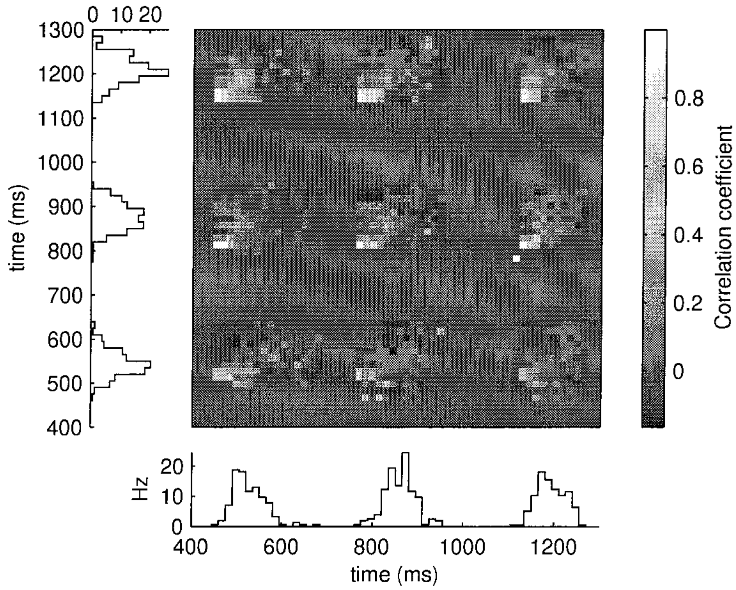

Figure 3.28: JPSTH of synthetic, Poisson-process-generated spike trains based on the PSTHs, excitabilities, and latencies found by the latency and excitability analysis. The analysis methods are described in chapter 2; the spike train generating methods are described in section 2.7.

ing estimated latencies, as found by the latency search method, is shown in Figure 3.30. Again, they vary on the scale of tens to hundreds of seconds.

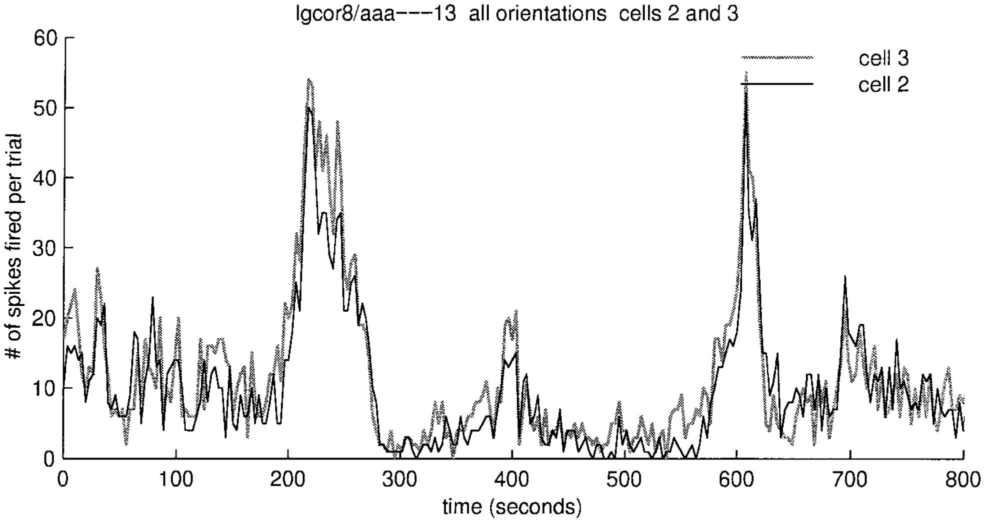

Figure 3.29: Total number of spikes fired in each trial as a function of trial start time. The curve for each cell is composed of 500 data points (500 trials).

Chapter 2 explained that covariations in mean firing rates (i.e., excitability correlations), such as those shown in Figure 3.29, are linked to positive DC values of the covariogram. A glance at the covariogram in the top panel of Figure 3.26 confirms that this holds true for this record (#lgcor8/aaa—13/133.75°): most of the covariogram is clearly above zero. Thus, record #lgcor8/aaa—13/133.75° is an example of how looking at the DC value of the correlogram can immediately allow one to tell whether there were covariations in the mean firing rates of the two cells. The central peak of the covariogram of this record, showing this tell-tale positive DC value, can also be found as Figure 2a in Sillito et al. (1994).

Aertsen et al. (1989) propose that "one may use the procedure of "slicing" the spike trains into different sections", each one with roughly similar mean firing rates, and "compare the results across sections" as a method for avoiding excitability correlation peaks. The methods of chapter 2 (in particular the *direct* estimation of the excitability correlation contribution to a correlogram) have been proposed here as more quantitative and complete ways to address this type of issue. Nevertheless, the result of the slicing procedure recommended by Aertsen et al. is shown below in Figure 3.31 for completeness. The result of applying the slicing technique shows that, after slicing, the remaining contributions to the covariogram are not statistically significant. Interest-

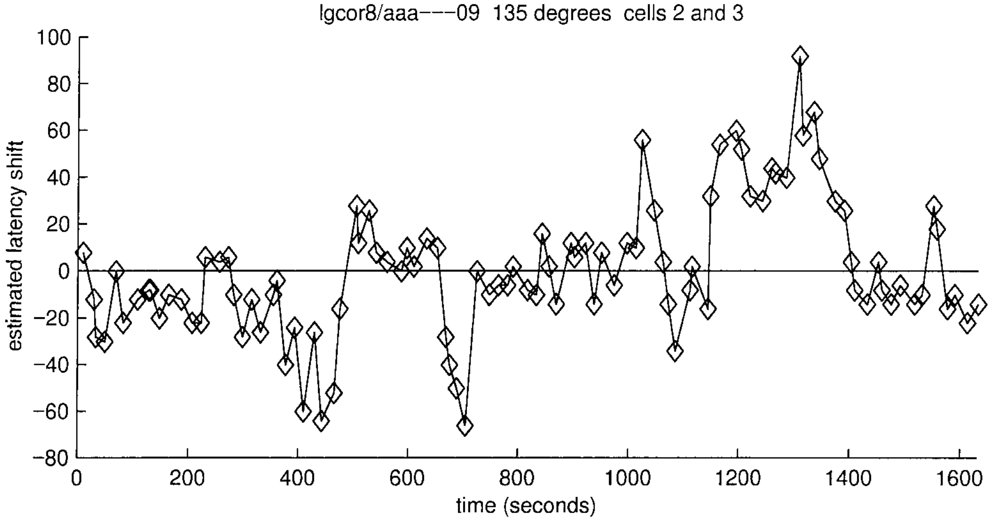

Figure 3.30: Latency estimate (as found by the latency search method), shown as a function of trial start time.

ingly, though not significant, a small peak near $t = 0$ remains. No side peaks are evident. It is possible that the remaining peak indicates the a weak but existing short time scale correlation, in addition to the longer time scale correlations the analysis has so far focussed on. If such a short time scale correlation component exists, however, its signal is too small to be statistically significant compared to the much stronger long time scale component. It is likely that in order to detect it, experiments where long time scale covariations were explicitly checked for and avoided, during the experiment itself, would need to be carried out.[16]

---

[16] This could be done by monitoring for stable mean firing rates and latencies, for example.

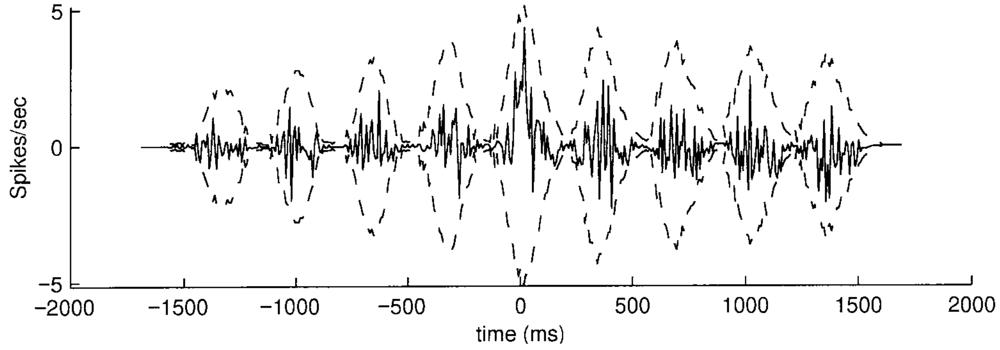

Figure 3.31: Result of slicing the record #lgcor8/aaa—13/133.75° into different sections, each with roughly constant mean firing rates. Each slice had at least three trials within it, and was normalised by the geometric mean of the number of spikes produced by each cell in that slice (Palm et al. (1988)). The covariograms of different slices were then averaged together. The variances of the null hypotheses corresponding to each covariogram were averaged and divided by the square root of the number of slices used to produce a final null hypothesis variance. The average covariogram is shown as a solid line; dashed lines are significance limits, defined as the 2σ limits obtained from the square root of the final null hypothesis variance. Interestingly, though not statistically significant, a peak near $t = 0$ can still be clearly seen.

3.6 The covarying resting membrane potentials hypothesis

The analysis presented in this chapter has argued that the observed peaks in the covariograms are due to slow, orientation-independent covariations in the latency and excitability of the cells' responses. A simple mechanism for generating such covariations in a simple integrate-and-fire neuron would be covariations in the resting membrane potential of the two cells involved: at low resting potentials, one would expect long latencies and low excitabilities, while at high resting potentials one would expect short latencies and high excitabilities.[17]. Such linked variations can be seen in the responses illustrated in Figure 3.25: in trials where the cells fire more spikes than in other trials, they also tend to fire sooner than in other trials, and this is true for both cells. However, LGN cells cannot always be approximated as simple integrate-and-fire neurons: when their resting potential is lowered enough, low-threshold Ca^{++} channels deinactivate, and the cells respond with bursts of spikes (Jahnsen and Llinas (1984a,b); see section 3.2.1). Thus, while for high ranges of resting potential variation, the cells would fire almost always in tonic mode (this

---

[17] I am indebted to John Hopfield for pointing out that variations in an integrate-and-fire cell's resting potential (or, equivalently, firing threshold) would lead to linked excitability and latency variations in the cell's spiking response, as described.

would also be true for cells with very few Ca^{++} channels), for other ranges of resting potential variations the cells might fire sometimes in tonic mode, and sometimes in burst mode; and for still other ranges (near the low end of the spectrum), the cells might fire always in burst mode.

All three of these possibilities are found in the experimental data. The covariogram in the top panel of Figure 3.26 is an example of a covariogram of cells in the *tonic* regime: neither cell fired any Ca^{++} bursts during this experiment. This is not shown in that Figure, but can be seen in its raster plots, shown in Figure 3.25. For these cells, the left panel of Figure 3.32 shows how increases in the total number of spikes fired are associated with decreases in latency, as one would expect if their resting membrane potentials were covarying, and they were well approximated by itnegrate-and-fire neurons. Being in the tonic mode, such an approximation may be adequate. The covariogram in the middle panel of Figure 3.26 is an example of a covariogram of cells in the *mixed* bursty/tonic regime: both cells fired bursts in about half the trials of the experiment (not shown). If the resting potentials of the two cells were covarying, one would expect both cells to fire tonically in the same trials, and to go into burst mode in the same trials. The right panel of Figure 3.32 shows how the tendency to fire in bursts indeed covaried in the two cells. Finally, the covariogram in the bottom panel of Figure 3.26 is an example of a covariogram of cells in the *burst* mode: both cells fired a burst in almost every modulation of almost every trial. Figures 3.33 and 3.34 show the rasters of these two cells. There are very clear slow covarying changes in the latency of burst firing, even though most bursts appear to be fairly similar.

The next chapter explores, through comptational modeling, the hypothesis that peaks in the covariograms are due to covariations in the cells' resting membrane potentials. All three firing regimes described here (tonic only, mixed, and burt only) will be discussed.

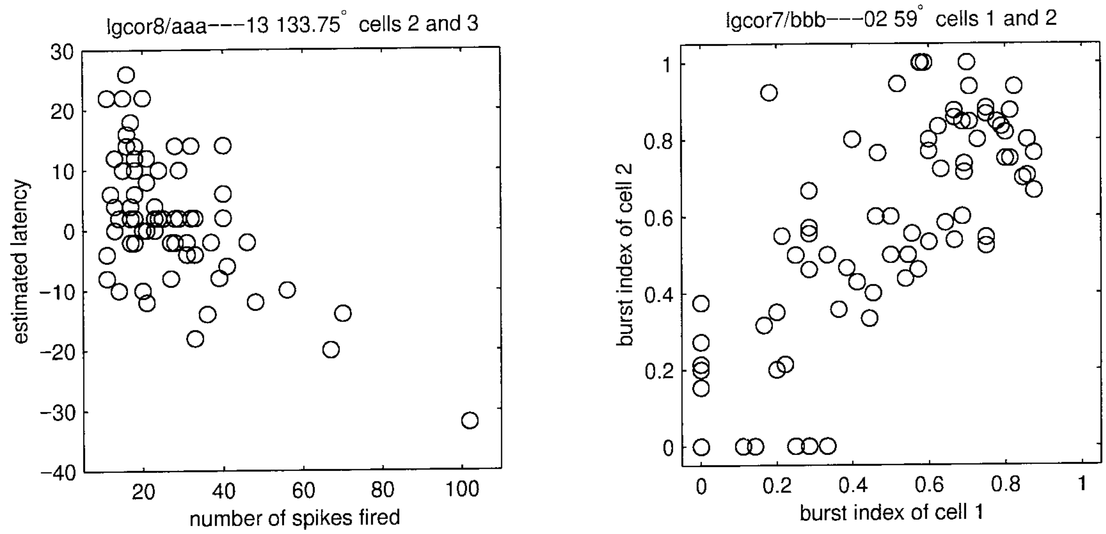

Figure 3.32: The left panel shows the estimated latency vs. the total number of spikes fired in each trial for the *tonic*-mode covariogram of Figure 3.26. The shortest latencies are associated with the greatest excitability. Correlation coefficient is -0.60. The right panel shows the burst index of cell 2 vs. the burst index of cell 1, for the data from the *mixed*-mode covariogram of Figure 3.26. The burst index is calculated for each raster for each cell as B_i =(# of spikes in Ca^{++} bursts)/(total number of spikes). Ca^{++} bursts were identified from the extracellular data using the criteria of Guido et al. (1992); Lu et al. (1992). The correlation coefficient of the burst index in the two cells is 0.86.

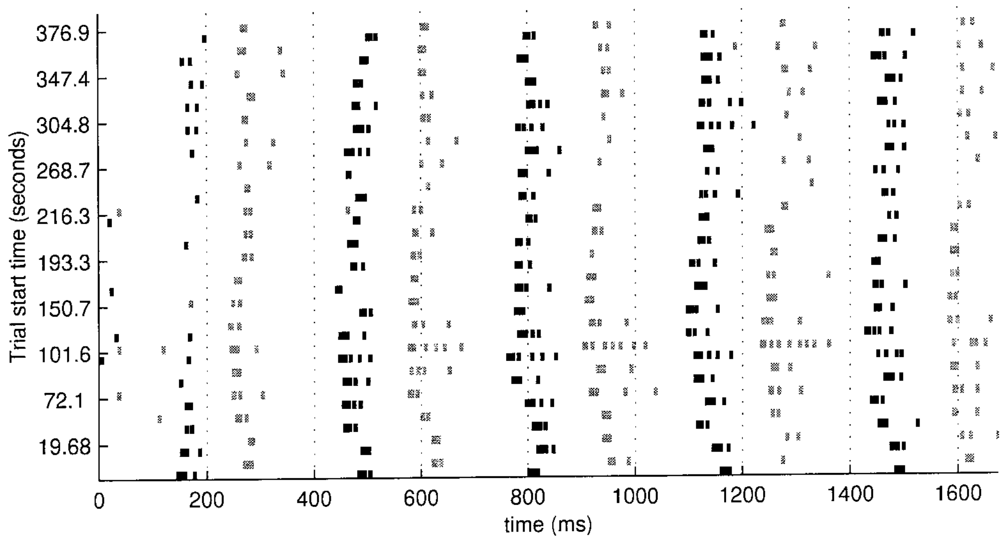

Figure 3.33: Rasters for the *burst*-mode covariogram of Figure 3.26. One cell is shown in black, the other in grey. Though most rasters fire similar numbers of spikes, slow changes in the latency of response can be clearly seen. This is most easily seen in Figure 3.34, a closeup of the rasters shown here. Data record lgcor9/ccc—11 at 90°.

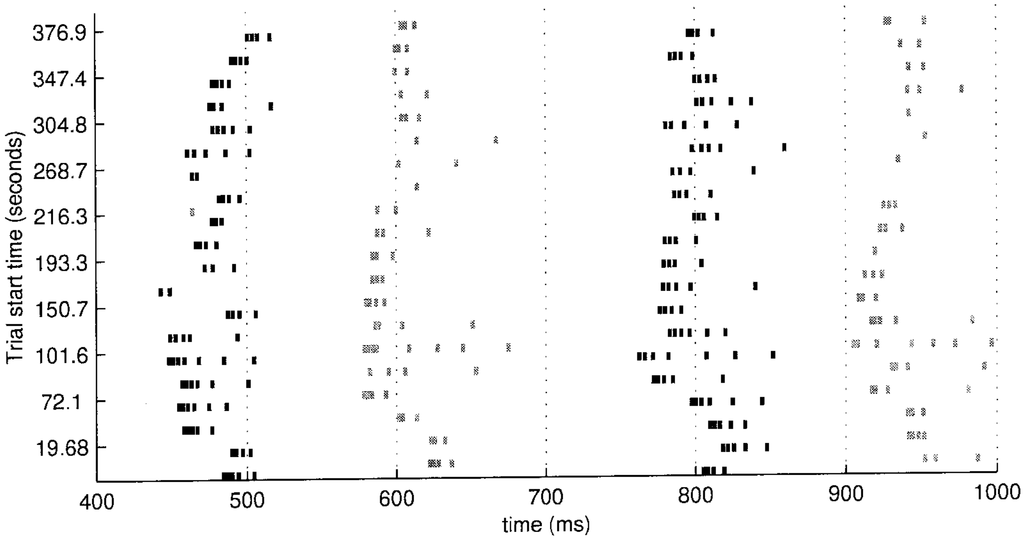

Figure 3.34: Closeup of rasters of Figure 3.33: time axis runs only from 400 to 1000 ms. Slow covariations in the latency of response can clearly be traced (e.g. responses to trials near start time 100 seconds had shorter latencies than responses to trials near start time 300 seconds. This is true for both cells, one of which is shown in black and the other in grey.). Trial start time is shown on the y-axis.

Chapter 4 Spike train correlations in the LGN of the cat: Modeling

4.1 Summary of the chapter

The analysis and observations described in the previous chapter led to the following hypothesis: the observed peaks in covariograms of LGN neuron pairs are caused by covariations in the resting membrane potentials (E_{rest}) of the two neurons. The modeling study to be described in this chapter explores this hypothesis.

Section 4.2 introduces the chapter by describing three different firing regimes seen in the data: *tonic* firing, *mixed* tonic and burst firing, and *burst* firing. The section explains how covarying resting membrane potentials could lead to covariogram peaks in each of these three regimes. Section 4.3 describes the model used to explore the hypothesis proposed here. The model is a slightly simplified version of Huguenard and McCormick's (1992) model of thalamic relay cells. Model responses to retinal ganglion cell (RGC) spiking input (spike times for which were drawn from data kindly provided by José-Manuel Alonso and R. Clay Reid) are illustrated for a variety of resting membrane potentials. The model clearly a dependence on E_{rest} of the latency and excitability, as well as the already well-known resting membrane potential dependence of burst vs. tonic firing response modes (Huguenard and McCormick, 1992). If the resting potentials of the two cells covary, their response properties will also covary, leading to a peak in their spike train covariograms. Section 4.4 describes the covariogram results in the three firing regimes. Covariations in E_{rest}, with the two cells remaining in the tonic regime, lead to broad covariogram peaks with positive DC values; covariations in E_{rest}, with the two cells remaining in the burst regime, lead to narrow peaks surrounded by slight inhibitory dips; and asymmetric covariations (where one cell stays in the tonic regime while the other varies from the tonic to the burst regime) lead to broad but asymmetric covariogram peaks. These types of peaks are all observed in the data. If the timescale of variations in the resting potentials is in the tens of seconds (as might be expected from known mechanisms which affect resting potential (McCormick and Vonkrosigk, 1992; Godwin et al., 1996; Murphy et al., 1994; Uhlrich et al., 1995; Sherman and Guillery, 1996)), then the side peaks seen in Figures such as 3.12 are also predicted by the "covarying resting membrane potentials" hypothesis.

4.2 Covariogram peaks from covarying resting potentials?

Three types of covarying firing patterns were seen in the experimental data. To each of these there corresponds a hypothesis as to how the covariogram peaks could be caused by covariations in the resting membrane potentials of the two neurons involved.

Case 1: Tonic regime. In this regime, the cells showed large covariations in the number of spikes produced in response to the visual stimulus, sometimes firing only 1 or 2 spikes per trial, sometimes as many as 50 (Figure 3.29). However, the cells never fired any bursts. Thus the T-type Ca^{++} channel appears to have been largely inactive. In the model, similar behaviour can be obtained by covarying E_{rest} for the two cells, but keeping the cells always in the tonic regime. As E_{rest} becomes more hyperpolarized, it will become more distant from the Na^{+} action potential firing threshold. Because of this, the LGN model cell will produce, in response to a given input, fewer spikes, and at a longer latency, than it would have at more depolarized values of E_{rest}. (see Figure 4.3). Thus, even while kept in the tonic regime, the response properties of LGN cells can covary due to covarying values of E_{rest}. This covariation is reflected in a positive peak in their spike train covariogram.

Case 2: Mixed burst/tonic regime. If the two cells covary from being in tonic response mode to being in burst response mode, then clearly their response properties will covary and a positive peak in their covariogram is to be expected. The change between burst and tonic response is a well-documented consequence of changes in E_{rest}. Hence this is another type of positive covariogram peak that can be caused by covariations in E_{rest}. Here, an asymmetric version of such a covariation will be considered: while the resting potential of the two cells will covary, the magnitude of the variation will be greater for one cell than the other. Thus one cell will vary within the tonic regime (in some trials firing many action potentials, in some very few) while the other will vary from firing tonically to firing in bursts. Since in this model the two cells are not interchangeable, the result will be an asymmetric covariogram peak.

Case 3: Burst regime. Since burst responses are composed of Na^{+} spikes which ride on top of a Ca^{++} action potential, they are fairly sterotyped: insofar as the Ca^{++} action potential is an all-or-none event (a rough approximation only) all bursts will be fairly similar, in the same way that all Na^{+} action potentials are fairly similar. But, as will be shown in the model below, the time at which the bursts happen, relative to the input, may be influenced by E_{rest}. Thus, in the burst regime, variations in E_{rest} lead to variations in latency, but with little variation in excitability.

Latency correlations generate peaks surrounded by inhibitory dips (see chapter 2). The width of the peaks will be largely determined by the width of the response bursts: thus, we expect narrow, 20-30 ms wide peaks (corresponding to the length of a typical burst), surrounded by inhibitory dips.

The geniculate relay cell model to be used must be capable of operating in both burst and tonic regimes. Following the strategy of Mukherjee and Kaplan (1995), a simplified version of McCormick and Huguenard's (1992) model was chosen as a relatively simple (having only a single compartment) model which nevertheless accurately reproduces *in vitro* both burst and tonic behaviour. In the model, two cells were separately simulated. On each trial, the resting membrane potential was set to be the same for both cells by controlling their leak K^+ conductances; this common resting potential varied slowly over different trials. (In a few runs, specifically those for the mixed burst/tonic regime, the two cells had covarying but not identical resting potentials.) Apart from the interaction through their resting potentials, the two simulations were independent and used independent inputs.

The leak K^+ conductance was chosen as the control of E_{rest} because there are known mechanisms, due both to signals from cortex (acting through metabotropic glutamate receptors) (McCormick and Vonkrosigk, 1992; Godwin et al., 1996), and to signals from the parabrachial region of the brain stem (acting through metabotropic acetylcholine receptors) (Murphy et al., 1994; Uhlrich et al., 1995; Sherman and Guillery, 1996), that can affect E_{rest} through the leak K^+ conductance.

4.3 Single-cell model

4.3.1 Description of the model

As stated above, the strategy of Mukherjee and Kaplan (1995) was followed and a simplified version of McCormick and Huguenard's very successful single-compartment model of geniculate cells (Huguenard and McCormick, 1992; McCormick and Huguenard, 1992) was used. The simplified version was chosen both for simplicity itself and for speed of simulation. Some runs were made with the full 10-channel model; the results in terms of the cross-correlogram shapes were essentially identical to those using the simplified model. The simplified model has a transient Na^+ and a delayed rectifier K^+ channel for action potential generation. Parameters for these were taken mostly from Belluzzi and Sacchi (1988, 1991). Details of the action potentials are not crucial for the results described here: they could have been replaced by a simple threshold-fire-and-reset mechanism. In contrast, the low-threshold Ca^{++} channel and its ensuing Ca^{++} spikes are very important for the results, since they determine the shape of burst responses and the resting-potential dependence of response mode. The kinetics and voltage-dependence of this Ca^{++} channel were taken from McCormick and Huguenard (1992). The channel was implemented as a permeability rather than a conductance, a Q_{10} of 3.55 was used for activation, and a Q_{10} of 2.8 was used for inactivation. Simulation "temperature" was 36° Celsius. The maximum permeability of the low-threshold Ca^{++} channel was slightly reduced from the full model to account for the elimination of several of the conductances of the full model, again following Mukherjee and Kaplan (1995). A leak Na^{+} and a leak K^{+} conductance were used in the model to set the membrane resting potential. On each trial, the two cells' membrane voltages were recorded, and upwards crossings above a -25 mV threshold were taken as spike initiation times. The resulting spike trains from many such trials were then used to compute covariograms.

The equations describing the model were numerically integrated using semi-implicit Backwards Euler integration with an adaptive timestep (Press et al., 1992). A single-compartment simulator program to do this was written in MATLAB 5. This platform was chosen over NEURON and GENESIS because of ease of integration of model results with MATLAB's data analysis capabilities, and because the model was simple enough that it was easily written despite not using a platform optimized for neural simulations.

4.3.2 Inputs to the model

The retinal ganglion cell (RGC) spike times used as input to the model cells were drawn from spike train recording data of an RGC in anesthetized paralysed cats, recorded during visual stimulation using drifting sine-wave gratings. This recording was made by José-Manuel Alonso and R. Clay Reid, who I thank for the kind sharing of their data. Like LGN cells, the response of RGCs to stimulation by drifting sine-wave gratings is similar to a halfwave-rectified sine wave (see Figure 3.1). The response can thus be divided into individual modulations, each corresponding to one period of the stimulus. Figure 4.1 shows rasters for 50 such modulations, together with their PSTH. The maximum firing rate of the RGC under these conditions was 144 Hz. The only ionotropic synaptic inputs to the two LGN cells in the model were those from RGCs. Each LGN cell in the model cell received input from a single RGC. Retinal ganglion cell spike times in response to approximately 3000 modulations of a drifting sine-wave grating were available to be used as inputs for the model. Different LGN cells and different trials or modulations in the model used as input different trials from the RGC spiking data. In the model, each incoming RGC spike generated a conductance change with an alpha-function shaped waveform:

$$g_{\mathrm{ampa}} = \hat{g}_{\mathrm{ampa}} \left[(t - t_0)/\tau\right] e^{-(t-t_0-\tau)/\tau} \quad \text{for} \quad t \geq t_0, \quad \text{zero otherwise} \tag{4.1}$$

$$I_{\mathrm{ampa}} = g_{\mathrm{ampa}} \, (E_{\mathrm{ampa}} - V_m). \tag{4.2}$$

Here t_0 is the arrival time of the spike, τ was set to 0.5 ms (for a conductance waveform with a timescale similar to that in Paulsen and Heggelund 1996), $\hat{g}_{\mathrm{ampa}}$ was the maximum conductance per spike, set to 27 nS (similar to Mukherjee and Kaplan 1995), and E_{ampa} was set to +40 mV. The effect of multiple incoming spikes on the conductance was additive. Since the time constant of the RGC AMPA inputs, 0.5 ms, is much shorter than the passive membrane time constant (which is on the order of 10 ms), these synaptic inputs essentially serve as charge dumpers, quickly raising the membrane voltage, which then slowly discharges through the membrane leak conductances. The decay time constant of these synaptic EPSPs is thus almost identical to the membrane time constant.

4.3.3 Responses of the model

A crucial feature of the model is its capacity to operate in either the burst or the tonic regime. Figure 4.2 follows Figure 11 of Mukherjee and Kaplan (1995) in illustrating how the responses of the model can change between these two regimes depending on the resting membrane potential. The same retinal spike train was used as input in both of the cases shown in the Figure, but for the top trace the resting membrane potential was set to -55 mV, and for the bottom trace it was set to -84 mV.

The responses in Figure 4.2 illustrate the responses with E_{rest} near the extremes of its possible values (E_K, the potassium reversal potential, was set to -95 mV, and the firing threshold for this model neuron is around -46 mV; these two are the limits of possible values of E_{rest}). How do the responses vary in between the two extremes? Panels A and C of Figure 4.3 show the spike train responses of the model to a particular, fixed, RGC input train, over a number of resting membrane potentials. For the purposes of the discussion below, the crucial points to notice in this Figure, seen particularly clearly in panel A, are that (*a*) there is a tonic regime ($\approx$ -52 to -70 mV), within

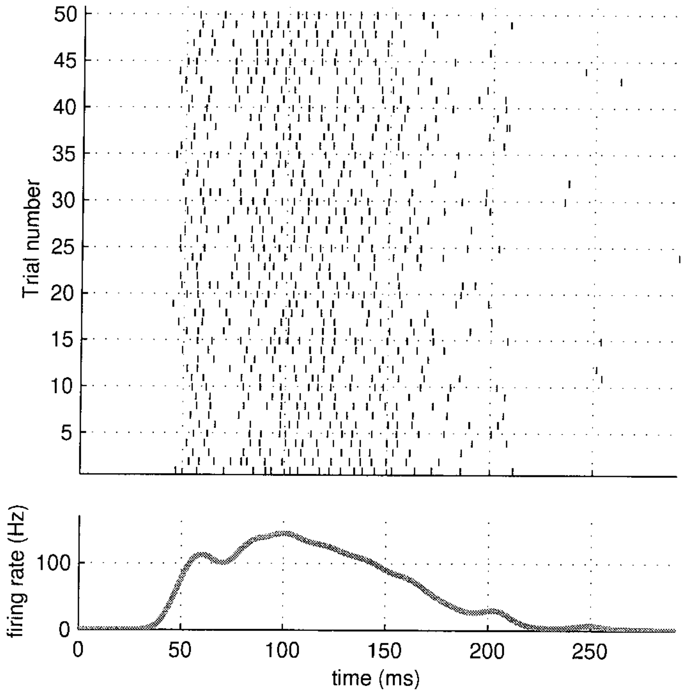

Figure 4.1: 50 different RGC responses to a modulation (i.e. one bar) of a drifting sine-wave grating. Data provided by José-Manuel Alonso and R. Clay Reid. The spike times will be used as input to the LGN model. The grey line on the bottom axes is the PSTH of 256 trials, smoothed by a Gaussian with standard deviation 6 ms.

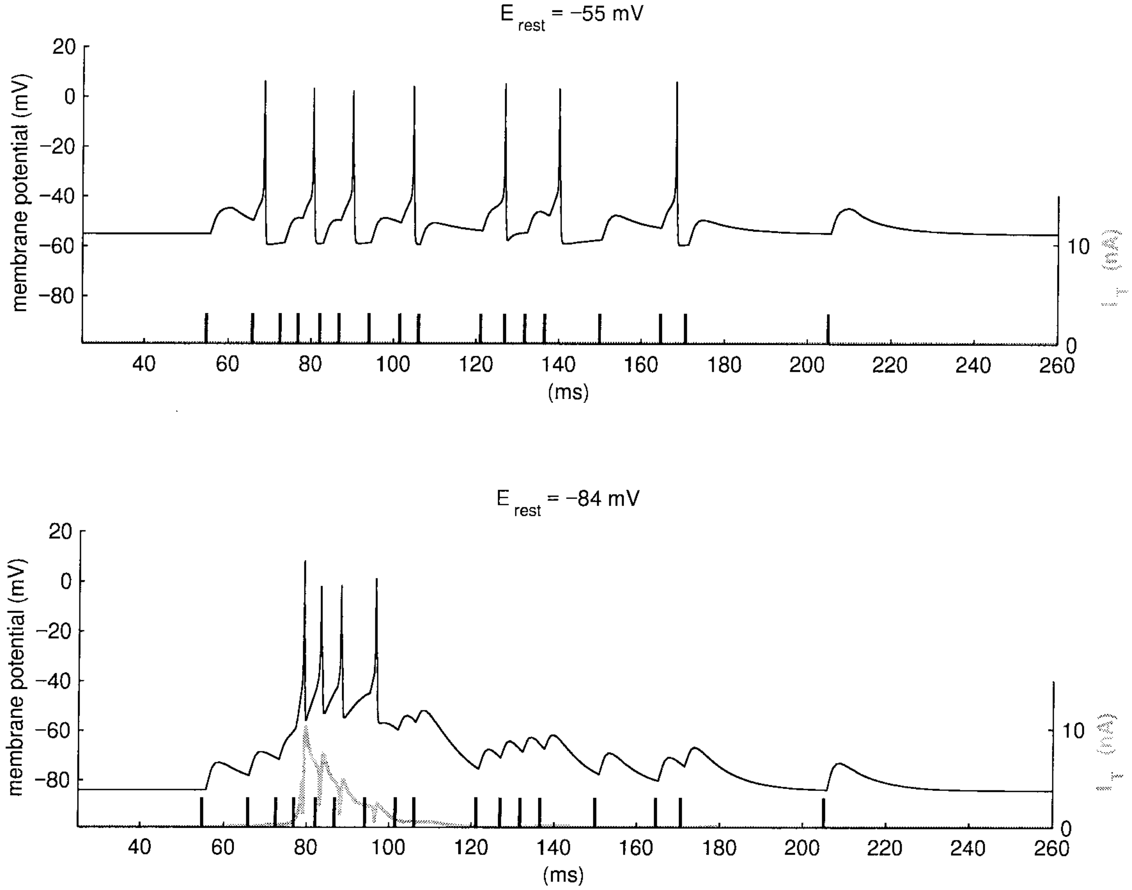

Figure 4.2: Responses of the model change from tonic to burst as a function of E_{rest}. The same RGC spike train input, shown as thick black vertical lines at the bottom of the two panels, was used in both simulation runs. The thin solid line is the membrane potential during the simulations (voltage scale at left). The thick grey line is the Ca^{++} current I_T (current scale at right) When the resting potential was -55 mV, there was no significant Ca^{++} current. This Figure follows Figure 11 of Mukherjee and Kaplan (1995).

which resting potential controls both response latency and "excitability" (i.e. number of spikes fired in response to a given input); (*b*) there is a burst regime ($\approx$ -78 to -87 mV), within which the resting potential controls mostly the latency of the stereotyped burst response; and (*c*) the resting potential controls the transitions between these regimes.

While Figure 4.3 showed the responses of a model LGN cell as a function of E_{rest}, for a fixed RGC input train, each of the panels of Figure 4.4 shows the responses as a function of different RGC input trains, for a fixed E_{rest}. This illustrates the variability in LGN spike output due to the variability in RGC spike input.

4.4 Modeling results

Two LGN model cells will now be used to explore the effect of slow covariations in resting membrane potentials on their spike train covariogram. The two cells' responses, over many trials of visual stimulation with drifting sine-wave gratings, were simulated. The resting membrane potentials of the two cells were kept fixed during each trial, but covaried slowly over different trials. In the model, as in the experiments, each trial lasted long enough for five bars of the drifting grating to bass over the receptive fields of the cells.

4.4.1 Tonic regime

In the tonic regime the dominant contribution to the covariogram peaks is the excitability correlation between the two cells, as indicated by the large positive DC values of the covariograms. Figure 4.5 compares experimental and model covariograms and JPSTHs. While the covariograms are fairly similar, the JPSTHs are quite different.

Figure 4.6 shows the model-produced and experimental rasters used to compute the covariograms and JPSTHs shown in Figure 4.5. An important difference between the two raster sets is that in the experimental data, the spike trains of each modulation all have their last spike at roughly the same time, regardless of the number of spikes fired; it is largely the leading edge of the response which changes timing depending on the total number of spikes fired. In contrast, in the model, it is the leading edge of the response that is relatively constant in its timing. This is due to the shape of the PSTH of the RGC input, which has a sharp leading edge, and is also due to the fact that the RGC used here often fired two closely spaced initial spikes; as illustrated in panel C of Figure 4.3, this often causes an LGN spike immediately after the RGC doublet. It is possible

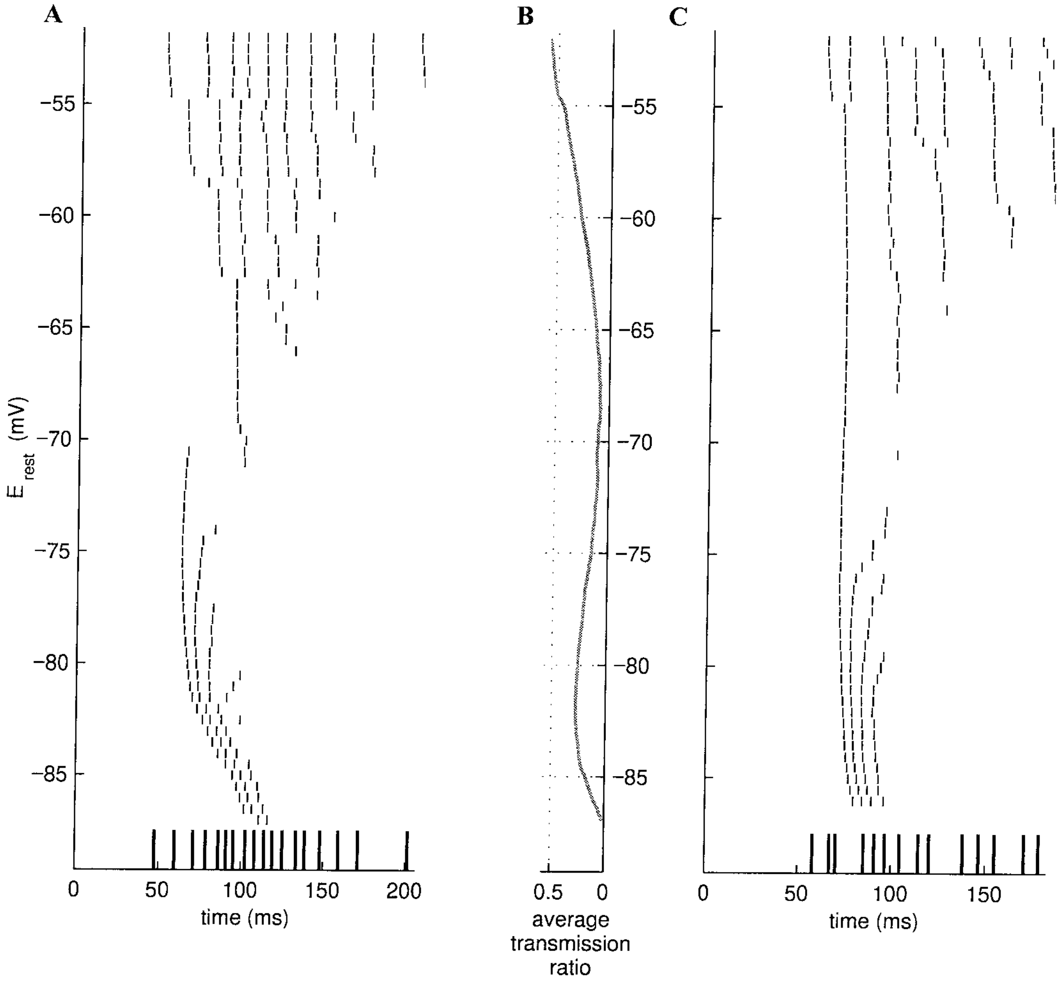

Figure 4.3: Responses of a model LGN cell as a function of resting membrane potential. Two different RGC spike train inputs, illustrated at the bottom of panels A and C as thick black vertical lines, were used; in each panel these inputs were kept constant. Panel B shows the transmission ratio (# of input spikes/# of output spikes), averaged over many different RGC input trains. Panel **A** illustrates how the resting potential, in the tonic regime ($\approx$ -70 to -52 mV), controls the number of spikes fired in response to the input; while in the burst regime ($\approx$ -87 to -75 mV), the effect of E_{rest} is mostly on the latency of the burst response. Panel **B** is an illustration of how two spikes, if spaced closely enough, can drive the cell to fire immediately after them at almost all values of E_{rest}.

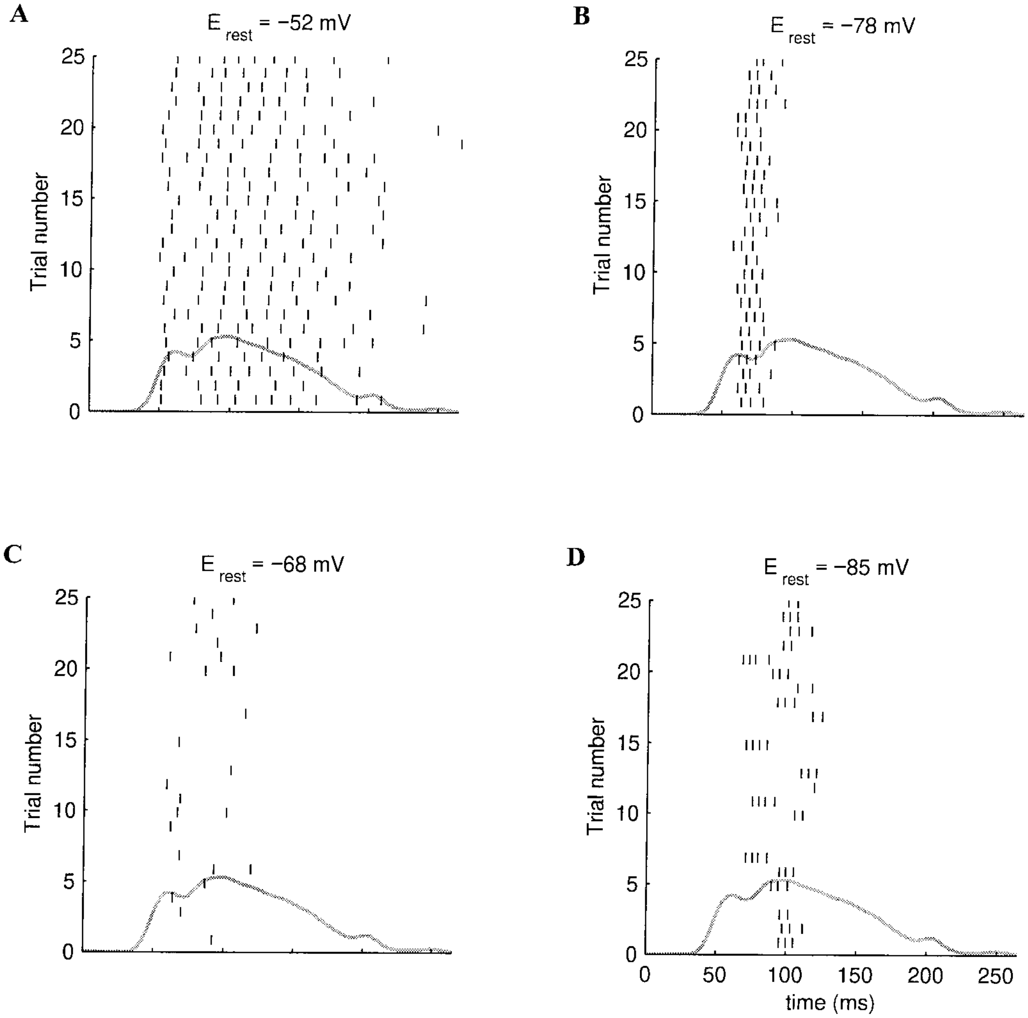

Figure 4.4: Responses of the model cell at four different values of E_{rest}. Each panel shows 25 runs using a different RGC input spike train on each run, and is thus an example of the variability in the LGN responses induced by the variability in RGC inputs. The thick grey line underlying the rasters is the PSTH of the RGC input; its maximum value is 140 Hz. The PSTH of LGN spikes (not shown) in panel A is close to a scaled version of the RGC's PSTH. This is in contrast to panels B, C, and D, where input and output PSTHs have very different shapes.

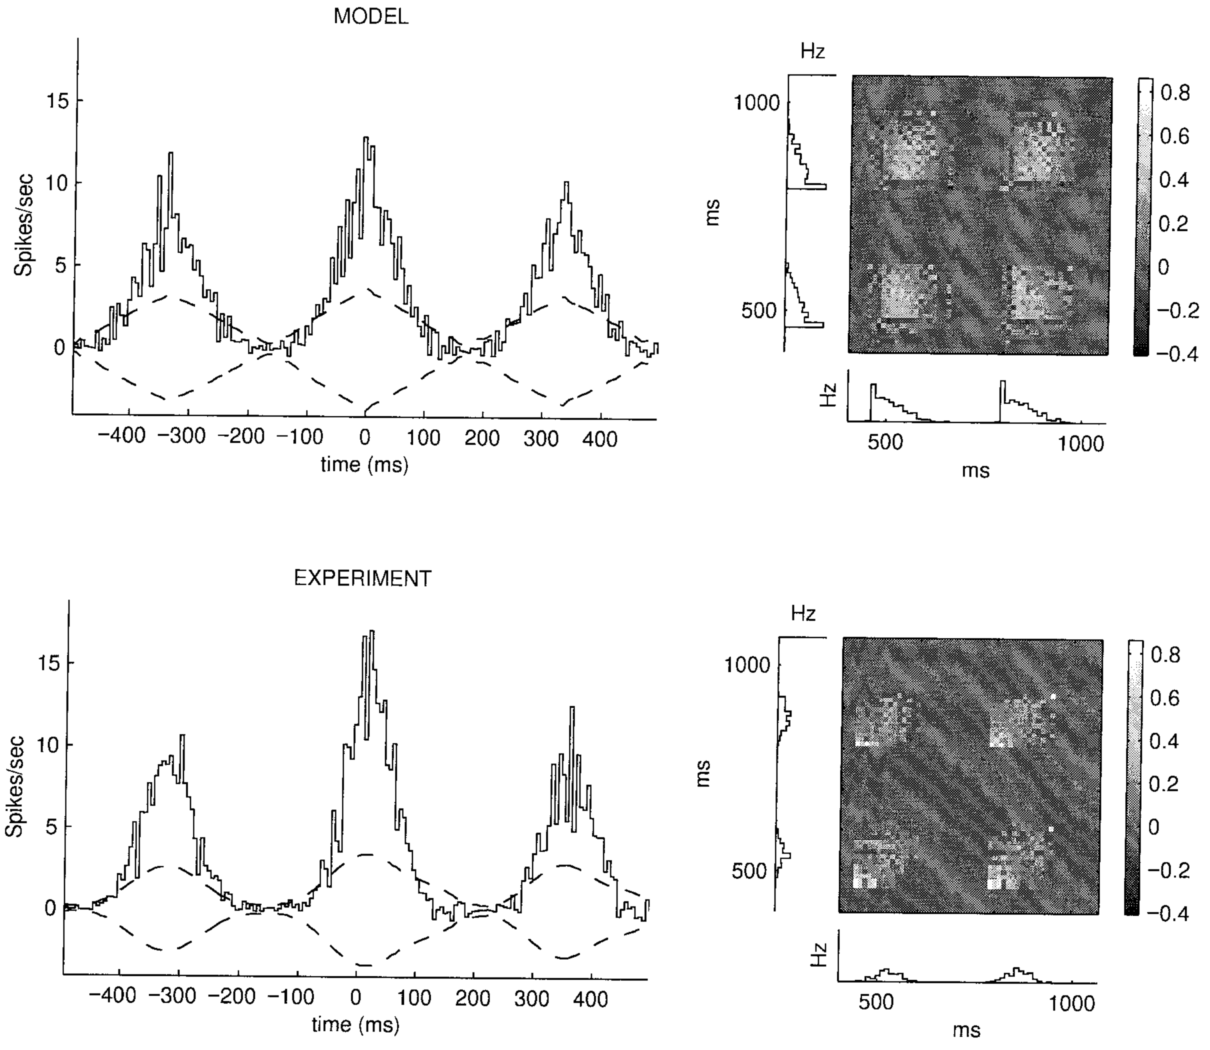

Figure 4.5: Tonic regime. Covariogram and JPSTH of experimental data (*top*) and model data (*bottom*). Binwidth for covariograms is 6 ms; for JPSTHs it is 12 ms. The axes for the PSTH next to each JPSTHs measure 90 Hz x 666 ms. The greyscale of the JPSTHs indicates correlation coefficient. Time in the JPSTH runs from left to right and from bottom to top (see chapter 2). Only a portion of the covariogram and JPSTHs, enough to see the central peak and the existence of the side peaks, are shown. While the two covariograms have similar features, the JPSTHs are clearly different.

that if a different RGC were used, different response properties (for example, a slow leading edge but a sharpt trailing edge fall-off), the model responses would match the experimental responses better and the JPSTHs would also match. The most salient feature of the experimental JPSTH, the positive leading peak (lower left-hand corner) is due to the covariation in the leading edge of the cells' response: as can be seen from the rasters, when one cell fires early, so does the other. The rasters also show that this feature varies only slowly with time, since responses to consecutive bars of the grating (e.g. response at 500 ms and at 800 ms in Figure 4.6) are correlated in the same way. The long time scale results in the side peaks of the covariogram and in the peaks in the off-diagonal modulations of the JPSTH. It is also possible that there is some spike adaptation process occurring in the real cells, responsible for the temporal asymmetry of their response properties (leading edge vs. trailing edge), which is not captured in the model.

4.4.2 Burst regime

Panel A of Figure 4.3 suggests that when the two LGN cells are operating mostly in the burst regime, covariations in the resting potential will lead mostly to covariations of the latency type. Figure 4.7 compares the covariogram and JPSTH of experimental data, in the burst regime, with the covariogram and JPSTH of model rasters, made with E_{rest} covariations within the burst regime. Since the timescale of variations of E_{rest} was greater than the temporal period of the simulated drifting grating stimulus, side peaks, similar to those in the experimental covariogram, are observed in the model data covariogram. The two covariograms have similar features: they have similar widths and are surrounded by similar inhibitory dips. The two JPSTHs also show similar features. Rasters of the spike times used in the computation of these covariograms and JPSTHs are shown in Figure 4.8. Clearly in this regime the model is capable of generating responses and covariograms very similar to those seen in the data.

4.4.3 Mixed burst/tonic firing regime

Figure 4.9 shows the covariograms and JPSTHs for model neurons and experimental neurons which fired both tonically and in bursts during the course of an experiment. Unlike the burst regime and the tonic regime, where the two model neurons always shared the same resting potential, here the resting potentials of the two model neurons covaried but were not identical. While neuron 1 (x axis on the JPSTH) always remained within the tonic regime, sometimes firing many

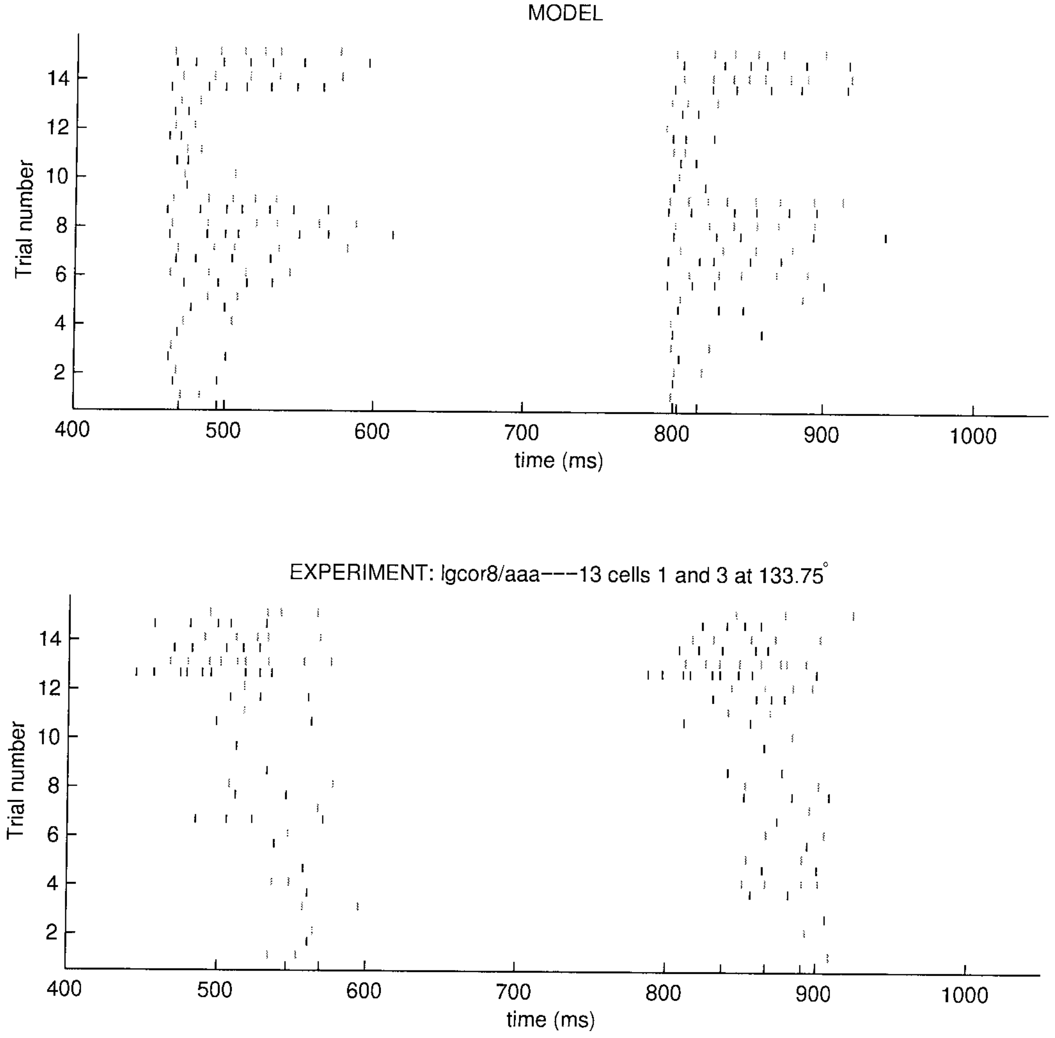

Figure 4.6: Rasters of the data used in Figure 4.5.

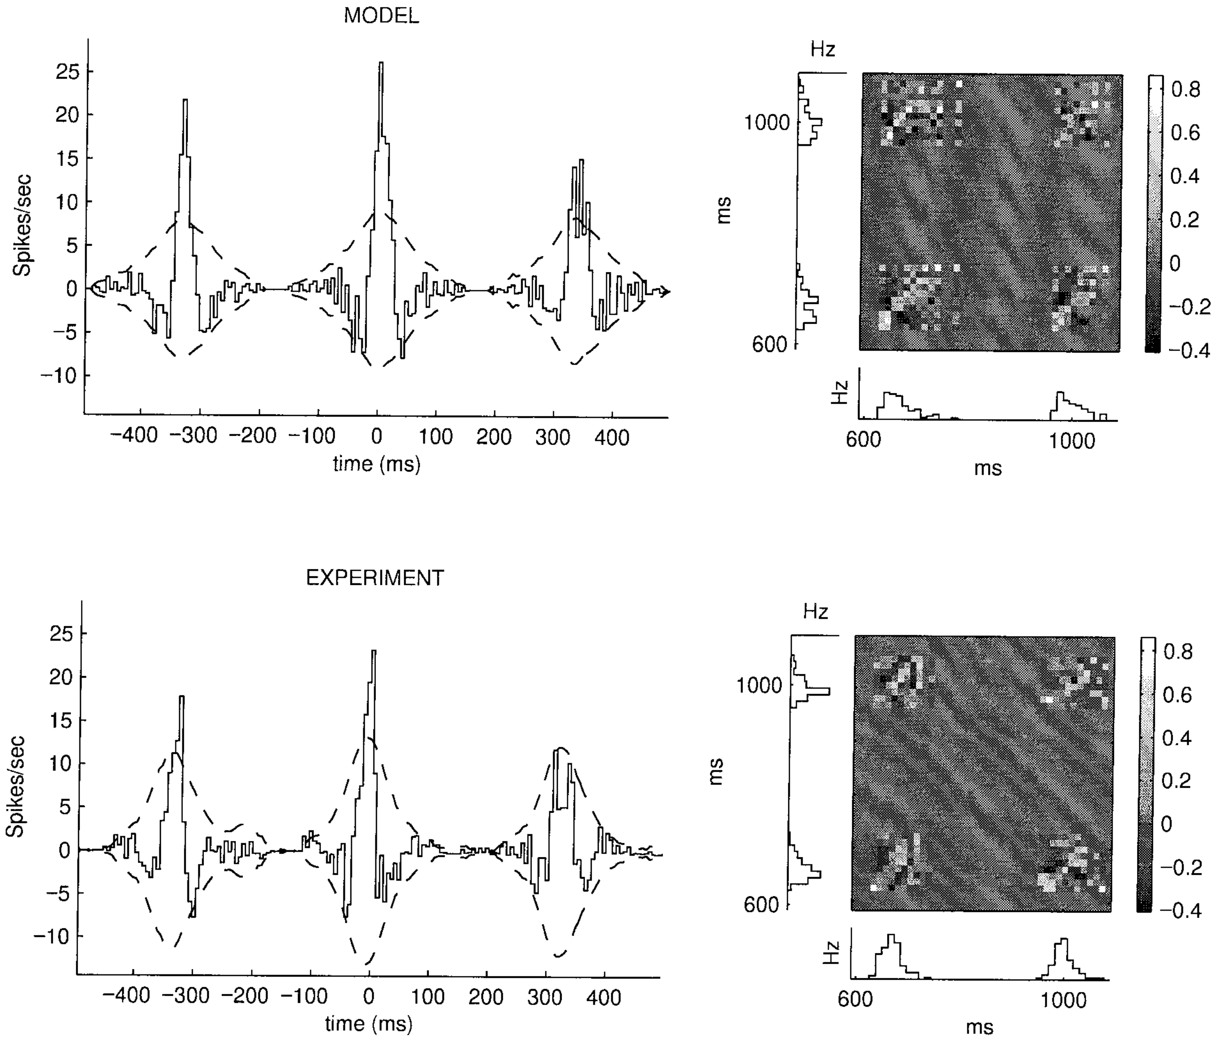

Figure 4.7: Burst regime. Covariogram and JPSTH of model data (*top*), and covariogram and JPSTH of experimental data (*bottom*). Binwidth for covariograms is 6 ms; for JPSTHs it is 12 ms. The axes for the PSTH next to each JPSTHs measure 130 Hz x 500 ms. The greyscale of the JPSTHs indicates correlation coefficient (see chapter 2). The model captures most of the features of the experimental covariogram and JPSTH.

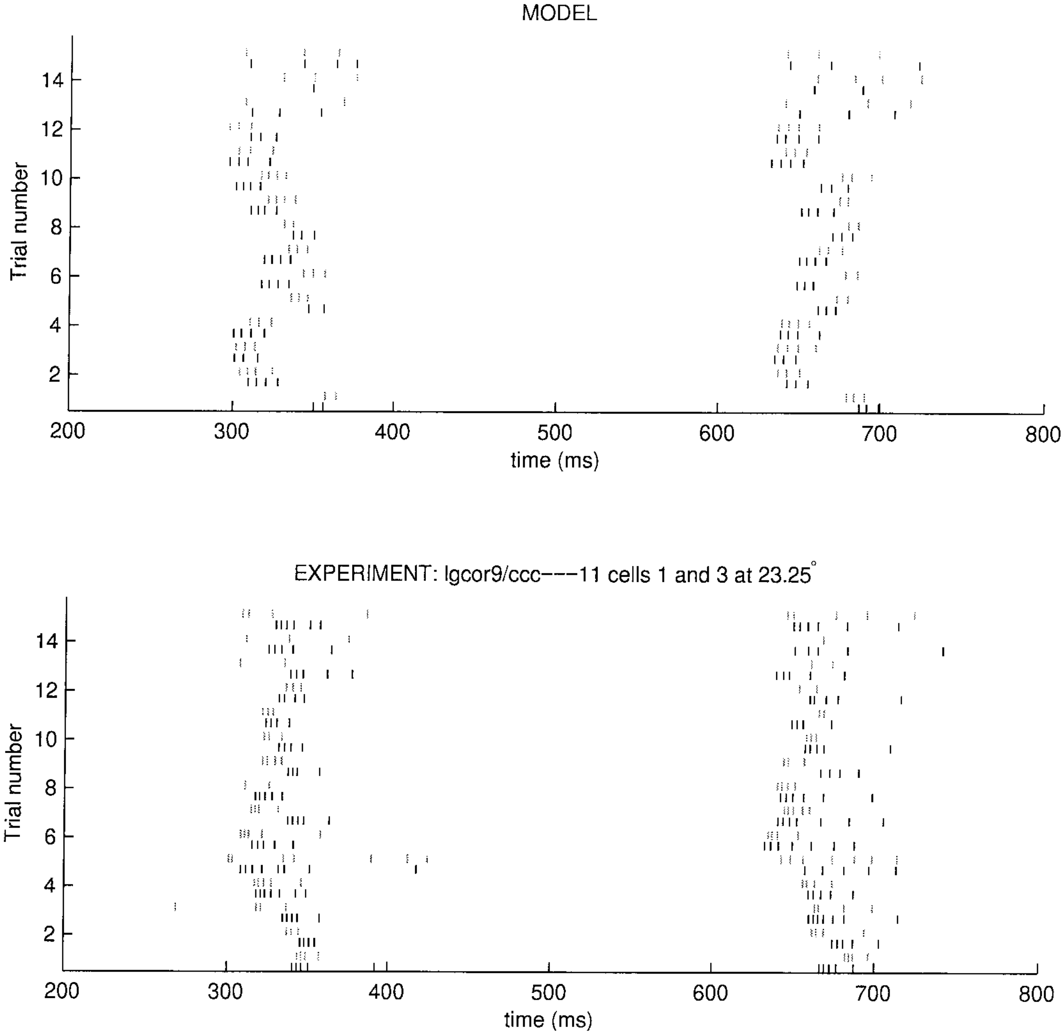

Figure 4.8: Rasters of the data used in Figure 4.7.

spikes, sometimes very few, neuron 2 (y axis on the JPSTH) sometimes fired many tonic spikes but sometimes fired in bursts. These kinds of trials were paired: when neuron 1 fired few spikes, neuron 2 responded with bursts, and when neuron 1 fired many tonic spikes, so did neuron 2. Since the response properties covary, the covariogram is not the zero function and has a peak. But since the two neurons are not interchangeable, the peak in the covariogram is asymmetric. The choice of resting potential asymmetry and consequent response types was guided by the responses seen in the experimental data; the asymmetric peaks can be seen in both the model data covariogram and in the experimental data covariogram: peaks have a slight dip to their left, and have a slightly slower fall-off to the right of the peak than to the left of the peak.

While the JPSTHs match well, the model data JPSTH sharing with the experimental one its very noticeable negative horizontal bands, the covariograms are clearly different. In particular, the DC value of the experimental covariogram is much greater than the DC value of the model data covariogram. Figure 4.10 shows experimental and model rasters. These suggest the origin of the difference between the covariograms: experimental bursts have fewer spikes than the model neuron's bursts. The model used here was the same as for the burst and tonic regimes: the only difference was the choice of resting potentials at which trials were run. Thus modelled bursts in the models of all three regimes were similar.[1] However, in the real system, different neurons will have different concentrations of low-threshold Ca^{++} channels, or differ in other characteristics, such that their bursts will not all be similar. If the model neuron in this regime produced fewer spikes per burst than those shown in Figure 4.10, the excitability correlations in the model covariogram of Figure 4.9 would be greater (since the burst trials of neuron 2 are those that are paired with the few-tonic-spikes trials of neuron 1). It is possible that under these conditions the experimental and model covariograms would match much more closely.

Conclusion of chapter

In conclusion, the very simple model used here can reproduce many characteristics of the experimental data and covariograms. A single model was used in all three regimes (tonic, burst, and mixed tonic/burst). The experimental data used for comparison with the model, however, came from three different pairs of neurons. Adapting the model to fit the characteristics of each of these cases would lead to closer fits to the covariograms. Even without adapting to these details, the general point made here is that covariogram shapes similar to those found experimentally,

---

[1] The burst response of the model was never used in the tonic regime, of course.

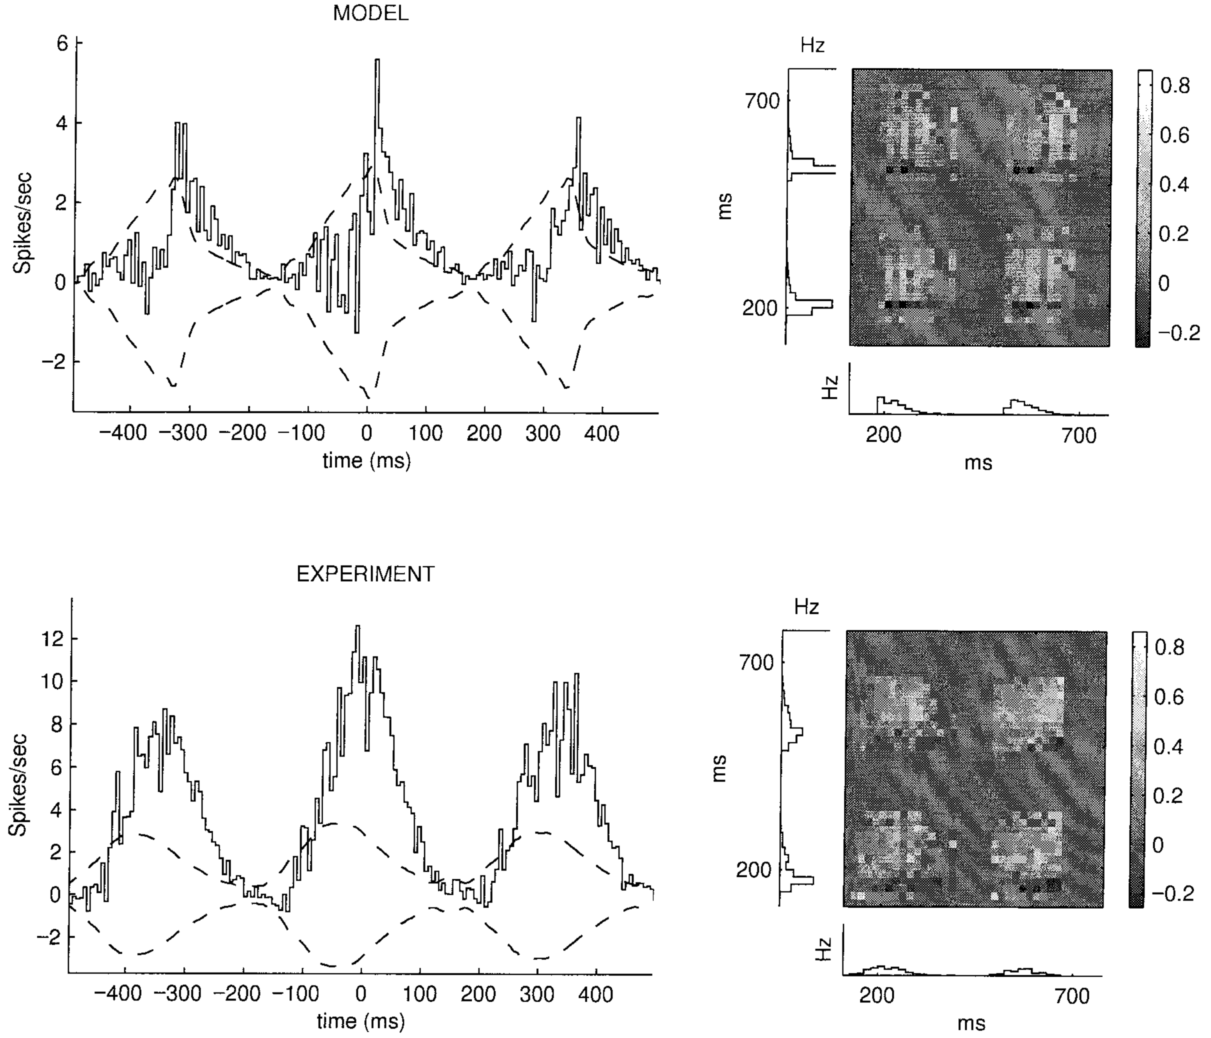

Figure 4.9: Mixed burst/tonic regime. Covariogram and JPSTH of model data (*top*), and covariogram and JPSTH of experimental data (*bottom*). Binwidth for covariograms is 6 ms; for JPSTHs it is 18 ms. The axes for the PSTH next to each JPSTHs measure 110 Hz x 666 ms. The greyscale of the JPSTHs indicates correlation coefficient (see chapter 2). The covariograms have a similar asymmetry (dips to the left of the peaks, a slower fall-off to the right of the peaks than to the left of the peaks) but the experimental one has clearly much stronger peaks and greater excitability covariations (this latter fact can be inferred from the DC values of the covariograms). The JPSTHs match well: the have the same type of negative horizontal bands (caused by the burstiness of one of the neurons– when neuron 2 (y axis) responds with a burst, neuron 1 (x axis) responds with fewer spikes than usual).

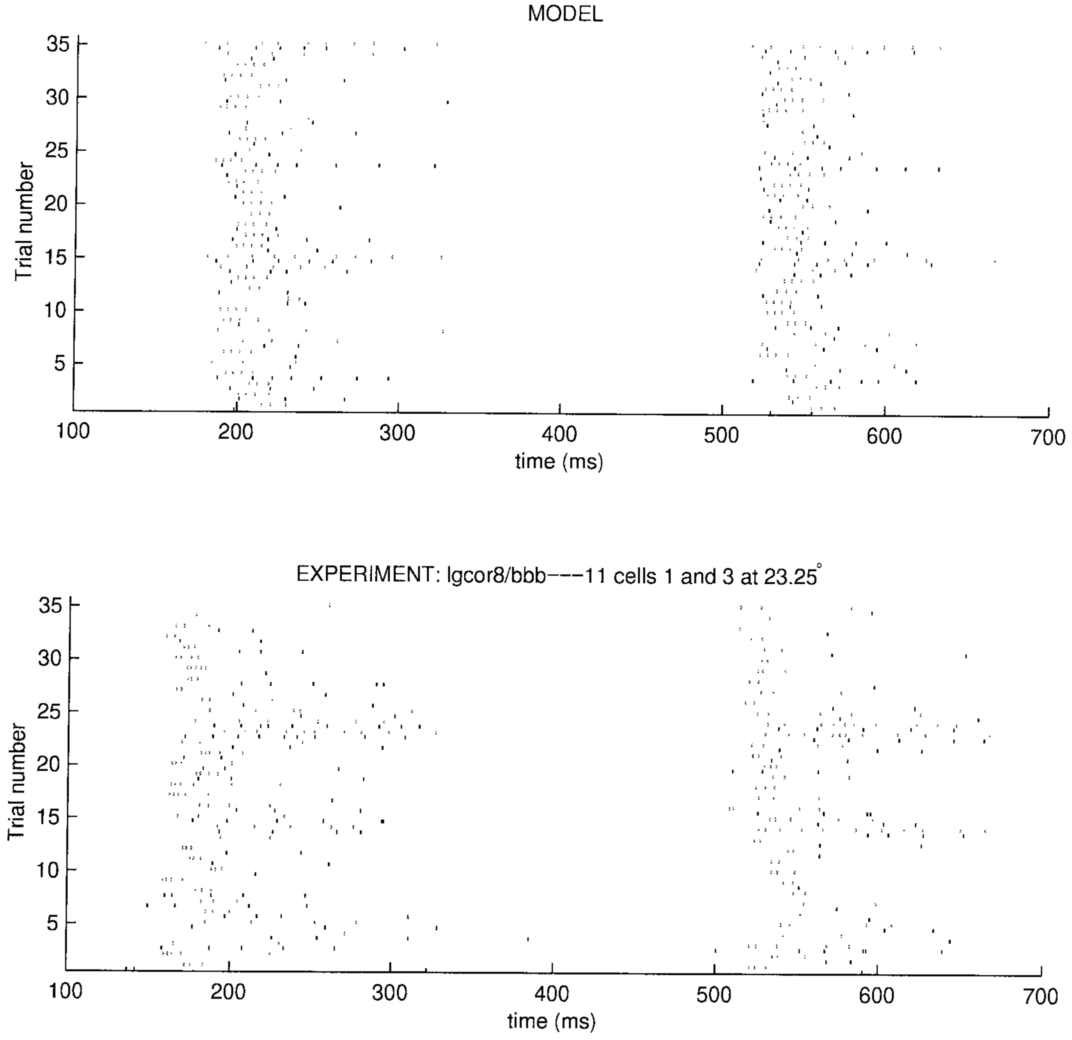

Figure 4.10: Experimental rasters (*top*) and model-produced rasters (*bottom*) in the mixed burst/tonic regime. The experimental data has fewer spikes per burst than the model.

including thin peaks (burst regime), broad peaks (mixed and tonic regimes), asymmetric peaks (mixed regime example), and side peaks of shape similar to the central peak (due to the slow timescale of the interaction), can be obtained from models where the only interaction between the two neurons is a slow, tens of seconds, resting potential covariation. Combined with the analysis of chapter 3, this supports the hypothesis that the dominant interactions between the neurons recorded by Sillito et al. were of the slow resting potential type and not due to short timescale spike timing synchronization.

Chapter 5 Discussion and Conclusions

This chapter will discuss and draw some conclusions from the results of chapters 3 and 4. The results of chapter 2 are discussed at the end of that chapter, in section 2.5.

Chapter 3 analysed the experimental data obtained by Sillito et al. (1994). Their experiment consisted of simultaneous extracellular recordings from pairs or trios of LGN neurons, with receptive fields between in 1° and 4° apart, recorded in anesthetized cats. When the visual system of the animals was stimulated with full-field drifting sine wave gratings, significant peaks in the covariograms of spike trains from pairs of neurons were seen. In their original report, Sillito et al. followed accepted practice and interpreted these peaks as evidence of spike timing correlations, or in their words, "synchronization of relay cell firing." They performed two key variations on the experiment: (1) When the stimulus was changed to small flashing spots centered over the receptive fields of the LGN neurons, the peaks in the covariograms were either abolished or strongly diminished, compared to the drifting gratings case.[1] While the flashing spots used in this variation of the experiment excite geniculate cells just as strongly as drifting gratings do, they excite primary visual cortex layer 6 cells (from where the corticogeniculate feedback projection originates) much less than drifting gratings do. (2) When drifting gratings were again used as stimuli, but visual cortex had been removed by aspiration, no cases of significant covariogram peaks were observed. These two observations seemed to indicate that the presence of the covariogram peaks was dependent on a strong feedback signal from visual cortex; furthermore, as stated above, the peaks were interpreted to mean that the feedback signal synchronized the relay cells' firing, thus increasing their effectiveness in driving their cortical cell targets (thanks to the temporal summation of their synchronized, and therefore over-lapping, EPSPs). Since it was thought that the synchronization of two geniculate relay cells occurred only when the stimulus driving them formed a continuous line between their two receptive fields (e.g. a bright bar of a grating, lying over both receptive fields, as opposed to two separate flashing spots, each one lying over one of the two geniculate receptive fields), this raised the intriguing and exciting possibility that the feedback from visual

---

[1] It is important to know that in the experiment the flashing squares trials were not interleaved with the drifting gratings trials; while data for the flashing squares was taken from the same cells from which drifting gratings data had previously been taken, this was done at a different time. Furthermore, the flashing squares generated geniculate PSTHs with a rather different time-course to those generated by drifting gratings.

cortex to the LGN might be involved in temporally "binding" together the representation of an object in the different cells of the visual system responding to it (Sillito et al. 1994; Singer 1994; for a review of the temporal binding hypothesis, see Singer and Gray 1995b).

Spike timing synchronization, as evidenced by covariogram peaks, is thought of as arising from common, coincident synaptic input (Moore et al., 1970). This type of interaction would not lead to covariogram peaks far from $t = 0$ on the covariogram time axis. Logically following this, Sillito et al. confined their covariogram analysis to a time window within ±200 ms of $t = 0$ (see figures in Sillito et al. 1994). But the full covariograms in fact show strong, significant peaks many 100s of ms away from $t = 0$ (see, for example, Figures 3.12 and 3.13). Furthermore, chapter 2 of this thesis demonstrated a hitherto underappreciated fact: covariogram peaks are not always due to spike timing correlations. They may be due to latency or excitability correlations; yet other possibilities also exist.[2] The analysis of section 3.5 shows that the covariogram peaks in Sillito et al.'s data can be well accounted for in terms of latency and excitability correlations, *without* a spike timing correlation component. These unexpected observations warranted a re-evaluation of the original analyses and interpretations. Sillito et al. generously shared their data for this effort.

Chapter 3 establishes that (1) the correlations have a time scale on the order of 10s of seconds (20–100 s); (2) the strength of the correlations is not significantly dependent on the orientation of the drifting grating used to stimulate the cells; and (3) responses to stimuli of a given orientation are not only correlated with responses to other trials using stimuli at the same orientation, but they are also correlated with responses to other trials using stimuli at a *different* orientation.

It must be stressed here that the analysis shows that the correlations have a long time scale, but it does not *prove* that there is not, in addition, a correlation at a shorter, perhaps milliseconds, time scale. It is possible that if there is such a short time scale correlation, it is masked by a partly noisy, much larger, slow time scale signal. The quality of the fits in the latency and excitability section (see section 3.5) suggests that if there is a short time scale correlation, it is weak: the latency/excitability predictions match the data extremely well (Figure 3.26), yet the predictions are based on parameters which are assumed to vary only every 4 or 5 seconds. (In the model on which the predictions are based, the covarying parameters are held fixed throughout every 2.5-second long trial, and trials are separated by 2 seconds). Other parameters (the PSTHs) have

---

[2] A passing reference to excitability correlations, as a confounder of spike timing correlation measurements, was made in Aertsen et al. (1989). Perkel et al. (1967) also mention excitability correlations, but merely to conclude that in the regime they were considering, they are unimportant. To my knowledge, latency correlations have not appeared anywhere in the literature.

faster variations, but these do not covary in the two cells and hence are not involved in interactions between them. Furthermore, the similarity of the peaks 100s of ms away from $t = 0$ to the peaks near $t = 0$ in the covariogram time axes suggest that the same mechanism is responsible for both of them. Finally, the peak height at short time scales is fully consistent with the peak heights found at time scales on the order of seconds (see Figure 3.17). All of these observations suggest that the dominant correlations in this data have a long time scale.

Since in the experiments different trials were separated by periods of no stimulation, the fact that the correlations last over several trials (section 3.4.2) indicates that the mechanism causing the correlations does not depend on brisk levels of activity in the visual system, and in particular does not depend on brisk levels of activity in V1, LGN, or the retina. These regions are known to fall relatively silent in the absence of visual stimulation.

The geniculate-projecting cells in layer 6 of visual cortex are tuned to the orientation of the visual stimulus. But the strength of the correlations does not depend significantly on the orientation of the stimulus (section 3.4.3). Therefore, if the mechanism causing the correlations is part of the corticothalamic pathway, it cannot depend on any particular arrangement between the location of the LGN receptive fields in visual space and the orientation tuning of any cortical cells contacting them. Cells tuned to all orientations would have to play a part in the correlations. One way this could occur would be for the mechanism involved in causing the correlations to be specific for each orientation, yet repeated for all orientations. For example, two correlated cells could both be contacted by a set of cortical cells spanning many different orientation tunings; if the correlations were caused by covariations in the efficacy of corticothalamic synapses, variations at synapses from cortical cells tuned to one orientation could be different to variations at synapses from cortical cells tuned to a different orientation. If each orientation set were covarying by the same magnitude in the two LGN cells being recorded from, then the peaks in the spike train covariograms of the two cells would be similar for all orientations, despite the variation at each orientation being separate from that at other orientations. However, section 3.4.4 shows that the responses at one orientation are correlated with the responses at different orientations. Thus the mechanism cannot be orientation-specific, even if equally strong for all orientations.

The latency and excitability analysis (section 3.5) of chapter 3 shows that the covariograms and JPSTHs found in the data can be well described in terms of covarying changes of only two global parameters (latency and excitability). Both latency and excitability are affected by resting membrane potential; this suggests that perhaps the covarying parameter underlying the correlated changes in latency and excitability in the two neurons is their resting membrane potential. In addition, several records show correlated changes in the burstiness of the two neurons being recorded from. This again could be due to correlated changes in the resting membrane potentials of the two neurons (Guido et al., 1992; Lu et al., 1992). Chapter 4 of the thesis explored this hypothesis by using a model of thalamic cells based on Hodgkin-Huxley-conductances. The model follows Mukherjee et al. (1995) in being a simplified version of Huguenard and McCormick's (1992) model. It can very accurately reproduce the experimental covariograms when the cells where in the burst mode, and it also reproduces (though less accurately), the covariogram peaks of cells in mixed burst/tonic firing mode and in tonic firing mode.

This set of results suggests that the key parameter covarying in the two cells, and causing the observed correlations, is resting membrane potential. The time scale on which this variable changes value is on the order of 20–100 seconds.

A number of inputs to LGN cells can affect its resting membrane potential (reviewed in Sherman and Koch 1986; McCormick 1992; Sherman and Guillery 1996; Sherman 1996), among them inputs from cortex (McCormick and Vonkrosigk, 1992; Godwin et al., 1996), and inputs from the para-brachial region of the brain stem (Uhlrich et al., 1995). The first have been shown to act through metabotropic glutamate receptors which have a time scale of tens of seconds (McCormick, 1992), and should thus be included as candidates for the mechanism causing the correlations. Inputs from the para-brachial region of the brain stem also act through metabotropic receptors and may have effects on similar time scales.

It has long been known that the level of burstiness of geniculate relay cells varies with the level of arousal of the experimental animal (McIlwain, 1972; Coenen and Vendrik, 1972; Kaplan et al., 1993). Burstiness becomes greater under deeper anesthesia. However, burst firing and changes in the level of burstiness have also been observed in fully awake cats (Mccarley et al., 1983; Guido and Weyand, 1995), suggesting that burst responses may play a role in visual function, perhaps by "alerting" cortex to the presence of particularly salient stimuli (Sherman 1996; see also Crick 1984). It is generally thought that brain stem inputs to the LGN are diffuse and global in nature; but whether the brain stem inputs project to selected cells found all over the LGN, or to *all* cells in the LGN has not been precisely established (Sherman and Guillery, 1996). Further, the corticothalamic inputs are known to be retinotopocally organized (Sherman and Koch, 1986), allowing the possibility of differential control in the resting membrane potential of different neurons in different parts of the LGN.

Unlike the preliminary report of Mukherjee et al. (1995), the data of Sillito et al. indicates that not all cells are correlated with each other at all times. None of the triple-cell recordings show covariograms (such as those exemplified in Figures 3.12 and 3.13) where all three pair combinations have strong peaks: in all recordings, one cell pair always covaries much more strongly than the other two, and in some records some cell pairs *anti*-covary. Further, Sillito et al. (1994) reported that no cases of correlations between cells of different polarity (On/Off) or type (X/Y) were found. Thus, the data of Sillito et al. holds evidence for in intriguing and hitherto undocumented phenomenon: that the control of burst vs. tonic firing, and control of excitability, is not uniform over the LGN but is subdivided into separate groups of cells. This control may be exerted via the corticogeniculate feedback projection. Whether the various groups are dynamic or static assemblies, what their properties are (are they simply the four possible combinations of On/Off/X/Y?), and what role differential control of each group plays in visual function are intriguing questions that deserve further investigation.

Bibliography

M. Abeles. *Corticonics*. Cambridge University Press, first edition, 1991.

A. M. H. J. Aertsen, G. L. Gerstein, M. K. Habib, and G. Palm. Dynamics of neuronal firing correlation - modulation of effective connectivity. *J. Neurophysiol.*, 61(5):900–917, 1989.

K. D. Alloway, M. J. Johnson, and M. B. Wallace. Thalamocortical interactions in the somatosensory system - interpretations of latency and cross-correlation analyses. *J. Neurophysiol.*, 70(3): 892–908, 1993.

J. M. Alonso, W. M. Usrey, and R. C. Reid. Precisely correlated firing in cells of the lateral geniculate nucleus. *Nature*, 383(6603):815–819, 1996.

O. Belluzzi and O. Sacchi. The interactions between potassium and sodium currents in generating action-potentials in the rat sympathetic neuron. *J. Physiol. London*, 397(MAR):127–147, 1988.

O. Belluzzi and O. Sacchi. A 5-conductance model of the action-potential in the rat sympathetic neuron. *Progress in Biophysics & Molecular Biol.*, 55(1):1–30, 1991.

C.D. Brody. Both NMDA and GABA_B synaptic mechanisms can account for corticofugally-induced cross-correlations in dLGN of the cat. *Society for Neuroscience Abstracts*, 22:1606, 1996.

C.D. Brody. Latency, excitability, and spike timing correlations. *Society for Neuroscience Abstracts*, 23:14, 1997.

V. A. Casagrande. A 3rd parallel visual pathway to primate area V1. *Trends in Neurosci.*, 17(7): 305–310, 1994.

A.M.L. Coenen and A.H.J. Vendrik. Determination of the transfer ratio of cat's geniculate neurons through quasi-intracellular recordings and the relation with the level of alertness. *Exp. Brain Res.*, 14:227–242, 1972.

F. Crick. Function of the thalamic reticular complex - the searchlight hypothesis. *Proc. Natl. Acad. Sci. USA Biological Sciences*, 81(14):4586–4590, 1984.

Y. Dan, J. J. Atick, and R. C. Reid. Efficient coding of natural scenes in the lateral geniculate nucleus - experimental test of a computational theory. *J. Neurosci.*, 16(10):3351–3362, 1996.

A. Das and C. D. Gilbert. Receptive-field expansion in adult visual cortex is linked to dynamic changes in strength of cortical connections. *J. Neurophysiol.*, 74(2):779–792, 1995.

C. Enroth-Cugell and J.G. Robson. The contrast sensitivity of retinal ganglion cells of the cat. *J. Physiol. London*, 187:517–552, 1966.

G. M. Ghose, I. Ohzawa, and R. D. Freeman. Receptive-field maps of correlated discharge between pairs of neurons in the cats visual cortex. *J. Neurophysiol.*, 71(1):330–346, 1994.

D. W. Godwin, J. W. Vaughan, and S. M. Sherman. Metabotropic glutamate receptors switch visual response mode of lateral geniculate nucleus cells from burst to tonic. *J. Neurophysiol.*, 76(3):1800–1816, 1996.

C. M. Gray and D. A. Mccormick. Chattering cells - superficial pyramidal neurons contributing to the generation of synchronous oscillations in the visual cortex. *Science*, 274(5284):109–113, 1996.

W. Guido, S. M. Lu, and S. M. Sherman. Relative contributions of burst and tonic responses to the receptive-field properties of lateral geniculate neurons in the cat. *J. Neurophysiol.*, 68(6): 2199–2211, 1992.

W. Guido and T. Weyand. Burst responses in thalamic relay cells of the awake behaving cat. *J. Neurophysiol.*, 74(4):1782–1786, 1995.

Y. Hata, T. Tsumoto, H. Sato, K. Hagihara, and H. Tamura. Development of local horizontal interactions in cat visual cortex studied by cross-correlation analysis. *J. Neurophysiol.*, 69(1): 40–56, 1993.

S. Hochstein and R. Shapley. Quantitative analysis of retinal ganglion cell classifications. *J. Physiol. London*, 262:237–264, 1976.

J. R. Huguenard and D. A. McCormick. Simulation of the currents involved in rhythmic oscillations in thalamic relay neurons. *J. Neurophysiol.*, 68(4):1373–1383, 1992.

H. Jahnsen and R. Llinas. Electrophysiological properties of guinea-pig thalamic neurons - an invitro study. *J. Physiol. London*, 349(APR):205–Continues, 1984a.

H. Jahnsen and R. Llinas. Ionic basis for the electroresponsiveness and oscillatory properties of guinea pig thalamic neurons invitro. *J. Physiol. London*, 349(APR):227–247, 1984b.

E. Kaplan, P. Mukherjee, and R. Shapley. *Contrast Sensitivity*, chapter Information Filtering in the Lateral geniculate Nucleus. MIT Press, 1993.

K. Kirkland and G.L. Gerstein. Models of cortically induced correlation observed in LGN neurons. *Society for Neuroscience Abstracts*, 22:906, 1996.

J. Kruger and F. Aiple. Multimicroelectrode investigation of monkey striate cortex - spike train correlations in the infragranular layers. *J. Neurophysiol.*, 60(2):798–828, 1988.

D. K. Lee, C. Koch, and J. Braun. Spatial vision thresholds in the near absence of attention. *Vision Res.*, 37(17):2409–2418, 1997.

S. M. Lu, W. Guido, and S. M. Sherman. Effects of membrane voltage on receptive-field properties of lateral geniculate neurons in the cat - contributions of the low-threshold Ca^{++} conductance. *J. Neurophysiol.*, 68(6):2185–2198, 1992.

R. W. McCarley, O. Benoit, and G. Barrionuevo. Lateral geniculate nucleus unitary discharge in sleep and waking - state-specific and rate-specific aspects. *J. Neurophysiol.*, 50(4):798–818, 1983.

D. A. McCormick. Neurotransmitter actions in the thalamus and cerebral-cortex and their role in neuromodulation of thalamocortical activity. *Progress in Neurobiol.*, 39(4):337–388, 1992.

D. A. McCormick and J. R. Huguenard. A model of the electrophysiological properties of thalamocortical relay neurons. *J. Neurophysiol.*, 68(4):1384–1400, 1992.

D. A. McCormick and M. von Krösigk. Corticothalamic activation modulates thalamic firing through glutamate metabotropic receptors. *Proc. Natl. Acad. Sci. USA*, 89(7):2774–2778, 1992.

J. T. McIlwain. Nonretinal influences on the lateral geniculate nucleus. *Invest. Ophtalmol. Visual Sci.*, 11:311–322, 1972.

G.P. Moore, J.P. Segundo, D.H. Perkel, and H. Levitan. Statistical signs of synaptic interaction in neurons. *Biophysical Journal*, 10:876–900, 1970.

P. Mukherjee and E. Kaplan. Dynamics of neurons in the cat lateral geniculate nucleus - in-vivo electrophysiology and computational modeling. *J. Neurophysiol.*, 74(3):1222–1243, 1995.

P. Mukherjee, T. Ozaki, and E. Kaplan. What controls the transfer of information through the LGN? *Society for Neuroscience Abstracts*, 21:657, 1995.

M. H. J. Munk, L. G. Nowak, J. I. Nelson, and J. Bullier. Structural basis of cortical synchronization .2. effects of cortical-lesions. *J. Neurophysiol.*, 74(6):2401–2414, 1995.

P. C. Murphy, D. J. Uhlrich, N. Tamamaki, and S. M. Sherman. Brain stem modulation of the response properties of cells in the cats perigeniculate nucleus. *Visual Neurosci.*, 11(4):781–791, 1994.

L. G. Nowak, M. H. J. Munk, J. I. Nelson, A. C. James, and J. Bullier. Structural basis of cortical synchronization .I. 3 types of interhemispheric coupling. *J. Neurophysiol.*, 74(6):2379–2400, 1995.

L. P. O'Keefe and M. A. Berkley. Binocular immobilization induced by paralysis of the extraocular-muscles of one eye - evidence for an interocular proprioceptive mechanism. *J. Neurophysiol.*, 66(6):2022–2033, 1991.

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

A. Papoulis. *Probability, Random Variables, and Stochastic Processes.* McGraw-Hill, second edition, 1984.

O. Paulsen and P. Heggelund. Quantal properties of spontaneous epscs in neurons of the guinea-pig dorsal lateral geniculate nucleus. *J. Physiol. London*, 496(3):759–772, 1996.

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

W. Press, S.A. Teukolsky, W.T. Vettering, and B.P. Flannery. *Numerical Recipes in C.* Cambridge University Press, second edition, 1992.

R.W. Rodieck, J.D. Pettigrew, P.O. Bishop, and T. Nikara. Residual eye movements in receptive-field studies of paralyzed cats. *Vision Research*, 7:107–110, 1967.

S. M. Sherman. Functional-organization of the W-cell, X-cell, and Y-cell pathways in the cat - a review and hypothesis. *Progress in Psychobiology and Physiological Psychol.*, 11:233–314, 1985.

S. M. Sherman. Dual response-modes in lateral geniculate neurons - mechanisms and functions. *Visual Neurosci.*, 13(2):205–213, 1996.

S. M. Sherman and R. W. Guillery. Functional-organization of thalamocortical relays. *J. Neurophysiol.*, 76(3):1367–1395, 1996.

S. M. Sherman and C. Koch. The control of retinogeniculate transmission in the mammalian lateral geniculate nucleus. *Experimental Brain Res.*, 63(1):1–20, 1986.

A. M. Sillito. Gaba mediated inhibitory processes in the function of the geniculo-striate system. *Progress in Brain Res.*, 90:349–384, 1992.

A. M. Sillito, J. Cudeiro, and P. C. Murphy. Orientation sensitive elements in the corticofugal influence on center-surround interactions in the dorsal lateral geniculate nucleus. *Experimental Brain Res.*, 93(1):6–16, 1993.

A. M. Sillito, H. E. Jones, G. L. Gerstein, and D. C. West. Feature-linked synchronization of thalamic relay cell firing induced by feedback from the visual cortex. *Nature*, 369(6480):479–482, 1994.

W. Singer. Neurobiology - a new job for the thalamus. *Nature*, 369(6480):444–445, 1994.

W. Singer and C. M. Gray. Visual feature integration and the temporal correlation hypothesis. *Ann. Rev. Neuro.*, 18:555–586, 1995a.

W. Singer and C. M. Gray. Visual feature integration and the temporal correlation hypothesis. *Ann. Rev. Neuro.*, 18:555–586, 1995b.

J. M. Sprague, P. Deweerd, E. Vandenbussche, and G. A. Orban. Orientation discrimination in the cat and its cortical loci. *Progress in Brain Res.*, 95:381–400, 1993.

J. P. Turner, N. Leresche, A. Guyon, I. Soltesz, and V. Crunelli. Sensory input and burst firing output of rat and cat thalamocortical cells - the role of NMDA and non-NMDA receptors. *J. Physiol. London*, 480:281–295, 1994.

D. J. Uhlrich, N. Tamamaki, P. C. Murphy, and S. M. Sherman. Effects of brain stem parabrachial activation on receptive-field properties of cells in the cats lateral geniculate nucleus. *J. Neurophysiol.*, 73(6):2428–2447, 1995.

R. Vogels and G. A. Orban. How well do response changes of striate neurons signal differences in orientation? - a study in the discriminating monkey. *J. Neurosci.*, 10(11):3543–3558, 1990.