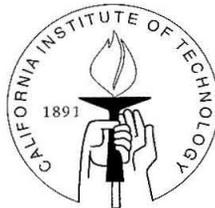


Analysis of Neuronal Dynamics in Behaving Animals

Thesis by
Bijan Pesaran

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



California Institute of Technology
Pasadena, California

2002

(Submitted 15 October, 2001)

© 2002

Bijan Pesaran

All Rights Reserved

Abstract

This thesis presents a set of mathematical techniques for analyzing neural activity and then applies them to data from a variety of experiments. Imaging data, in which movies of brain activity are recorded, is considered first and ways to suppress noise and characterize the signal are explored. Data from the cortex of behaving monkeys is then considered. These techniques are used to analyze the activity of populations of neurons during eye movements. Oscillations are found that encode information predicting "where" and "when" an eye movement is made. Activity in other parts of the brain involved in reaching and the perception of motion are then analyzed and shown to encode other information in a similar way. These results show neuronal dynamics may be used by the brain to process information during behavior.

Acknowledgements

A great many people contributed to my education during the writing of this thesis. First and foremost among them is my advisor, Partha Mitra. He has given me the tools to do research for which I will always be grateful. I am also indebted to Richard Andersen for the opportunities he has provided me with. His lab is an exciting place to work. I would also like to thank the other members of my committee: Mike Cross and Ken Libbrecht. I hope this thesis gives a flavour of how techniques from the physical sciences can contribute to understanding a biological system.

During the early stages of my thesis I was fortunate enough to spend time at Bell Labs. Thanks go to Peter Littlewood for making this possible. I would also like to thank David Tank, David Kleinfeld, Michale Fee, Sebastian Seung, Dan Lee, and Winfred Denk for my invaluable experience there. Special thanks go to David Thomson for his development of the multitaper method for spectral analysis.

At Caltech, I am grateful to have worked with the many members of the Andersen Lab including Boris Breznen, Chris Buneo, David Dubowitz, Betty Gillikan, Bradley Greger, Murray Jarvis, Daniella Meeker, Cierina Reyes, and Viktor Shcherbatyuk.

This thesis is due to the generosity of those whose data is analyzed in these pages including Aaron Batista, Ken Britten, Chris Buneo, Larry Cohen, Kerry Delaney, Alan Gelperin, X. Hu, David Kleinfeld, Rodolfo Llinas, Bill Newsome, Seiji Ogawa, John Pezaris, Jim Prechtl, Urs Ribary, Maneesh Sahani, Mike Shadlen, Larry Snyder, Kamil Ugurbil and the countless others that make their work possible. I hope that my intention to do experiments is evidence of how much I value their work.

I have wanted to do research for as long as I can remember. For this, I thank my parents.

I dedicate this thesis to my wife, Tina, whose patience shows me understanding, whose tenderness gives me strength, and whose love is forever.

Contents

| | |
|---|------------|
| Abstract | iii |
| Acknowledgements | iv |
| Introduction | 1 |
| 1 Data analysis techniques | 5 |
| 1.1 Time series analysis techniques | 5 |
| 1.1.1 Time domain versus frequency domain: resolution and nonstationarity | 6 |
| 1.1.2 Digitization rate, Nyquist frequency, Fourier transforms | 8 |
| 1.1.3 Conventional spectral analysis | 9 |
| 1.1.4 Autoregressive spectral estimation | 11 |
| 1.1.5 Multitaper spectral analysis | 12 |
| 1.1.6 Choice of the bandwidth parameter | 20 |
| 1.2 Analysis of multivariate data | 20 |
| 1.2.1 Eigenmode analysis: SVD | 21 |
| 1.2.2 Space-frequency SVD | 25 |
| 1.3 Local frequency ensemble and jackknife error bars | 27 |
| 1.4 Spectral analysis of point processes | 30 |
| 1.4.1 Counting process | 30 |
| 1.4.2 Interval process | 32 |
| 1.4.3 Poisson process | 32 |
| 1.4.4 Renewal process | 33 |
| 1.5 Spectral analysis of hybrid processes | 33 |

| | | |
|----------|--|-----------|
| 2 | Analysis of dynamic brain imaging data | 35 |
| 2.1 | Introduction | 35 |
| 2.2 | Different brain imaging techniques | 36 |
| 2.2.1 | Imaging techniques and their spatiotemporal resolution | 36 |
| 2.2.2 | Sources of noise | 37 |
| 2.3 | Description of data sets | 38 |
| 2.4 | Analysis techniques for different modalities | 39 |
| 2.4.1 | Magnetoencephalography | 40 |
| 2.4.1.1 | Preprocessing | 40 |
| 2.4.1.2 | Time-frequency analysis | 41 |
| 2.4.1.3 | Multichannel spectral analysis | 45 |
| 2.4.2 | Optical imaging | 48 |
| 2.4.2.1 | Spectral analysis of PC time series | 48 |
| 2.4.2.2 | Removal of physiological artifacts | 48 |
| 2.4.2.3 | Space-frequency SVD | 49 |
| 2.4.3 | Magnetic resonance imaging | 52 |
| 2.4.3.1 | Removal of physiological artifacts | 53 |
| 2.4.3.2 | Space-frequency SVD | 55 |
| 2.5 | Discussion | 56 |
| 3 | Temporal structure during working memory in Macaque parietal cortex | 60 |
| 3.1 | Introduction | 60 |
| 3.2 | Materials and methods | 62 |
| 3.2.1 | Animal preparation | 62 |
| 3.2.2 | Behavioral task | 62 |
| 3.2.3 | Electrophysiological recordings | 63 |
| 3.2.4 | Spike sorting | 65 |
| 3.2.5 | Data analysis | 66 |
| 3.3 | Results | 70 |

| | | |
|----------|--|------------|
| 3.3.1 | Mean spike and LFP responses | 70 |
| 3.3.2 | Temporal structure in spike activity | 72 |
| 3.3.3 | Temporal structure in LFP activity | 77 |
| 3.3.4 | Coherency between spiking and LFP activity | 84 |
| 3.4 | Discussion | 90 |
| 3.4.1 | Understanding the temporal structure in spike activity | 90 |
| 3.4.2 | Temporal structure in the LFP | 90 |
| 3.4.3 | Coherent gamma band activity during working memory | 91 |
| 4 | Interspike interval correlations in cortical activity: Data analysis and modeling | 93 |
| 4.1 | Introduction | 93 |
| 4.2 | Methods | 94 |
| 4.2.1 | Behavioral paradigms | 94 |
| 4.2.2 | Experimental recordings | 95 |
| 4.2.3 | Data analysis | 96 |
| 4.3 | Results | 97 |
| 4.3.1 | Stochastic process models | 97 |
| 4.3.2 | Interval analysis of cortical activity | 101 |
| 4.4 | Discussion | 106 |
| | Appendix A: Visual stimuli induce waves of electrical activity in turtle cortex | 107 |
| A.1 | Methods | 108 |
| A.1.1 | Preparation | 108 |
| A.1.2 | Optical recording | 109 |
| A.1.3 | Electrophysiology | 109 |
| A.1.4 | Stimuli | 110 |
| A.1.5 | Data reduction | 110 |
| A.2 | Results | 113 |

| | | |
|---|--|------------|
| A.2.1 | Basic response | 113 |
| A.2.2 | Large-scale dynamics | 116 |
| A.3 | Discussion | 120 |
| Appendix B: The role of nonlinear dynamics of the syrinx in the vocalizations of a songbird. | | 121 |
| B.1 | Introduction | 121 |
| B.2 | Methods | 121 |
| B.2.1 | Isolated syrinx | 121 |
| B.2.2 | Syrinx response to parameter changes | 122 |
| B.2.3 | Stroboscopic imaging | 123 |
| B.2.4 | Spectral analysis | 123 |
| B.3 | Results | 123 |
| B.3.1 | Analysis of bird song | 123 |
| B.3.2 | Mechanism of song generation | 126 |
| B.3.3 | Syringeal modeling | 129 |
| B.4 | Discussion | 129 |
| Appendix C: List of abbreviations | | 133 |
| Bibliography | | 134 |

List of Figures

| | | |
|------|--|----|
| 1.1 | Sample continuous time series | 10 |
| 1.2 | Autoregressive spectral estimate | 10 |
| 1.3 | Slepian functions: Time domain | 12 |
| 1.4 | Slepian functions: Frequency domain | 13 |
| 1.5 | Multitaper spectral estimate | 16 |
| 1.6 | F-statistic for artifact suppression | 18 |
| 1.7 | Singular value spectrum | 22 |
| 1.8 | Principal component spectra | 24 |
| 1.9 | Multivariate coherence spectrum | 27 |
| 1.10 | Space-frequency spatial eigenmodes | 28 |
| 1.11 | Jackknife error bars | 29 |
| 2.1 | Principal component time series of MEG data | 41 |
| 2.2 | 60 Hz artifact suppression | 42 |
| 2.3 | Prewhitening MEG data | 43 |
| 2.4 | Time-frequency spectrum of MEG data | 44 |
| 2.5 | Cross-frequency correlation matrix | 45 |
| 2.6 | Coherence map of MEG data | 46 |
| 2.7 | 35–45 Hz spatial eigenmodes of MEG data | 47 |
| 2.8 | SV coherence spectrum of limax data | 50 |
| 2.9 | Spatial eigenmodes of limax data | 51 |
| 2.10 | Phase gradients in spatial eigenmode of Limax data | 52 |
| 2.11 | Moving sinusoidal model of physiological artifacts | 54 |
| 2.12 | Estimated components of suppressed physiological artifacts | 55 |
| 2.13 | SV coherence spectrum of fMRI data | 56 |

| | | |
|------|---|-----|
| 2.14 | Spatial eigenmodes of fMRI data | 57 |
| 3.1 | The memory-saccade task | 63 |
| 3.2 | Broad-band extracellular potential | 64 |
| 3.3 | Spike and LFP mean responses | 71 |
| 3.4 | Area LIP spike spectrum | 72 |
| 3.5 | Spike population spectrum | 74 |
| 3.6 | Spike spectrograms | 76 |
| 3.7 | LFP spectrum at a single site | 78 |
| 3.8 | LFP spectrum population data | 79 |
| 3.9 | Spike and LFP tuning curves | 81 |
| 3.10 | LFP spectrograms | 82 |
| 3.11 | LFP power in gamma and beta bands | 83 |
| 3.12 | Spike-triggered average potential | 85 |
| 3.13 | Spike-field coherence at a single site | 86 |
| 3.14 | Spike-field coherence population data | 87 |
| 3.15 | Spike-field time-frequency coherency | 88 |
| 3.16 | Spike-field phase histogram | 89 |
| 4.1 | Simulated Poisson process | 98 |
| 4.2 | Simulated renewal process | 99 |
| 4.3 | Simulated renewal process with bursting | 100 |
| 4.4 | Simulated Markov renewal process | 100 |
| 4.5 | Sample LIP cell | 101 |
| 4.6 | Three sample MT cells | 103 |
| 4.7 | MT cell population average | 104 |
| 4.8 | Sample PRR cell | 105 |
| A.1 | Experimental procedure | 114 |
| A.2 | The spatiotemporal optical signal | 115 |
| A.3 | Spatial coherence of optical signal | 117 |

| | | |
|-----|---|-----|
| A.4 | Temporal evolution of the band-limited electrical activity. | 119 |
| B.1 | Spectral and temporal characteristics of zebra finch song | 125 |
| B.2 | Patterns of acoustic signals generated in the in-vitro syrinx | 128 |
| B.3 | Stroboscopic observation of in-vitro syringeal oscillations | 130 |
| B.4 | Model of airflow through a compliant constriction | 131 |

Introduction

The brain constitutes a complex dynamical system with a large number of degrees of freedom, so multichannel measurements are necessary to gain a detailed understanding of its behavior. Such multichannel measurements, made available by current instrumentation, include multielectrode recordings, optical brain images using intrinsic (Blasdel and Salama, 1986; Grinvald et al., 1992) or extrinsic (Davila et al., 1973) contrast agents, functional magnetic resonance imaging (fMRI) (Ogawa et al., 1992; Kwong et al., 1992) and magnetoencephalography (MEG) (Hamalainen et al., 1993). As a result of improvements in the capabilities of the measuring apparatus, as well as growth in computational power and storage capacity, the data sets generated by these experiments are increasingly large and more complex. The analysis and visualization of such multichannel data is an important piece of the associated research program, and is the subject of this thesis.

There are several common problems associated with the different types of multichannel data enumerated above. First, preprocessing is necessary to remove nuisance components, arising from both instrumental and physiological sources, from the data. Second, an appropriate representation of the data for purposes of analysis and visualization is necessary. Third, there is the task of extracting any underlying simplicities from the signal, mostly in the absence of strong models for the dynamics of the relevant parts of the brain. If there are simple features that are hidden in the complexity of the data, then the analytical methodology should be such as to reveal such features efficiently.

With the current exponential growth in computational power and storage capacity, it is increasingly possible to perform the above steps in a semiautomated way, and even in real time. In fact, this is almost a prerequisite to the success of multichannel measurements, since the large dimensionality of the data sets effectively precludes exhaustive manual inspection by the human experimenter. An additional challenge is to perform the above steps as far as possible in real time, thus allowing quick feedback into the experiment. Note that even

given the increases in storage capacity it is desirable to have ways of compressing the data while retaining the appropriate information, so as to prevent saturation of the available storage.

Problems such as the above are clearly not unique to neuroscience. Automated analysis plays an important role in the emergent discipline of computational molecular biology. Despite the current relevance of these problems, the appropriate analytical and computational tools are in an early stage of development. In addition, investigators in the field are sometimes unaware of the appropriate modern signal processing tools. Since little is understood about the detailed workings of the brain, a straightforward exploratory approach using crude analysis protocols is usually favored. However, given the increasing availability of computational resources, this unnecessarily limits the degree of knowledge that can be gained from the data, and at worst can lead to erroneous conclusions, for example when statistical methods are applied inappropriately (cf. Cleveland (1993) p.177). Alternatively, a superficial application of complex signal processing or statistical techniques can lead to results that are difficult to interpret.

An aspect of the data in question that cannot be emphasized enough is the fact that the data constitute time series, mostly multivariate. While techniques for treating static high-dimensional data are widely known and appreciated both in multivariate statistics and in the field of pattern recognition, the techniques for treating time series data are less well developed, except in special cases. Special consideration is given to analyzing the data in reference to the associated behavior of the animal.

The first chapter reviews some of the relevant analytical techniques for multivariate time series. In particular, a description of multitaper spectral methods (Thomson, 1982; Thomson and Chave, 1991; Percival and Walden, 1993). This is a framework for performing spectral analysis of univariate and multivariate time series that has particular advantages for the data at hand. A central issue is to be able to deal with very short data segments and still obtain statistically well behaved estimators. Reasons for this are that gathering long time series may be expensive (for example in fMRI), and that the presence of nonstationarity in the data due to behavior makes it preferable to use a short, moving analysis window.

Multitaper methods are particularly powerful for performing spectral analysis of short data segments. This chapter also presents the spectral analysis of point processes which can be treated in the same fashion as continuous time series.

The second chapter treats in succession dynamic brain imaging data gathered in MEG, optical, and fMRI experiments. Based on the analysis of actual data sets, protocols to remove artifactual components as well as to determine the structure of the signal are discussed. The properties of the data as multivariate time series are emphasized. Principal concerns are twofold: (1) characterization and removal of the typical artifacts, and (2) characterization of underlying structure of the signal left after removal of artifacts. Together with parts of chapter one, this work was published in the *Biophysical Journal* (Mitra and Pesaran, 1999).

The third chapter presents spiking and local field potential (LFP) activity from area LIP in the parietal lobe during a memory-saccade task. Using multitaper spectral analysis presented in Chapter 1 elevated power in the gamma band (25–90 Hz) in spiking and LFP activity is found during the memory period but not simple fixation. This activity is spatially tuned to the direction of the saccade. Spiking and LFP activity are also coherent in the gamma band but not low frequency. These results are evidence for memory fields defined by temporal structure that we term *dynamic memory fields* in which spiking is correlated with extracellular currents that give rise to gamma band LFP power. LFP activity in the beta frequency band (15–25 Hz) is related to preparation, and movement execution is also found. Therefore, temporal structure in the LFP in parietal cortex predicts both the direction and time of a planned movement which could accelerate the development of a cortical neural prosthesis.

The fourth and final chapter compares and contrasts spiking activity from three areas that contain gamma band activity, the lateral intra-parietal area (LIP), the parietal reach region (PRR), and the middle temporal area (MT), during three different tasks, the memory-saccade, the memory-reach and a visual motion discrimination task, respectively. Previous analysis of these data are extended by estimating interspike interval (ISI) correlations using the interval spectrum. This sheds new light on the nature of gamma band activity in

parietal cortex, and implications for stochastic process models for the data are discussed.

The thesis also contains two appendices. The first appendix presents work using optical imaging with an extrinsic contrast agent to study waves of electrical activity in the visual areas of turtle cortex. This work was published in the Proceedings of the National Academy, USA (Prechtl et al., 1997). The second appendix presents work studying the nonlinear dynamics of the syrinx in a song bird. This work was published in Nature (Fee et al., 1998). The final appendix contains a list of abbreviations.

Chapter 1 Data analysis techniques

1.1 Time series analysis techniques

Here we briefly review some methods used to analyze time series data. The aim here is not to provide a complete list of the relevant techniques, but discuss those methods which are directly relevant for the work to follow. In particular, in the next section, we provide a review of multitaper spectral analysis techniques (Thomson, 1982; Percival and Walden, 1993).

The basic example to be considered is the power spectral analysis of a single (scalar) time series, or an output scalar time series given an input scalar time series. The presence of temporal structure in neural activity is of much interest (Singer and Gray, 1995). Correlation functions are often estimated in the time domain to detect temporal structure but have long been known to suffer serious problems of estimation bias and variance which are exacerbated in the context of a behaving animal (for a review see Jarvis and Mitra (2001)). These problems cannot be addressed by pooling observations across a large period of time with different associated behaviors during the experiment as this leads to violation of the stationarity assumption and potential misinterpretation of the data. The relevant analysis techniques can be generally categorized under two different attributes: linear or nonlinear, and parametric or nonparametric. We will mostly be concerned with multitaper spectral techniques, for which the attributes are linear and nonparametric. Although we categorize the techniques as linear, note that spectra are quadratic functions of the data. Also, some of the spectral quantities we consider are other nonlinear functions of the data. We prefer nonparametric spectral techniques (e.g., multitaper spectral estimates) over parametric ones (e.g., autoregressive spectral estimates, also known as maximum entropy spectral estimates, or linear predictive spectral estimates). Some weaknesses of parametric methods in the present context are lack of robustness and lack of sufficient flexibility to fit data with

complex spectral content. The reader is referred to the literature for a comparison between parametric and nonparametric spectral methods (Thomson, 1982; Percival and Walden, 1993), and we also discuss this issue further below.

Methods based on time lag or delay embeddings will not be used as they characterize neurobiological time series as outputs of underlying nonlinear dynamical systems. These methods work if the underlying dynamical system is low-dimensional and if one can obtain large volumes of data so as to enable construction of the attractor in phase space. The amount of data needed grows exponentially with the dimension of the underlying attractor. On the one hand, it is true that neurobiological time series are outputs of rather nonlinear dynamical systems. However, in most cases it is not clear that the constraint of low dimensionality is met, except perhaps for very small networks of neurons. In cases where the dynamics may appear low-dimensional for a short time, nonstationarity is a serious issue, and precludes acquisition of very long stretches of data. One might think that nonstationarity could be accounted for by simply including more dynamical degrees of freedom. However, that also would require the acquisition of exponentially larger data sets. We constrain ourselves to spectral analysis techniques as opposed to the techniques indicated above. The reason for this is twofold: first, spectral analysis remains a fundamental component of the repertoire of tools applied to these problems, and as far as we are concerned, the appropriate spectral techniques have not been sufficiently well studied or utilized in the present context. Second, it remains debatable whether much progress has been made in understanding the systems involved using the nonlinear techniques (Theiler and Rapp, 1996; Rapp et al., 1994).

1.1.1 Time domain versus frequency domain: resolution and nonstationarity

In the neurobiological context, data are often characterized in terms of appropriate correlation functions. This is equivalent to computing corresponding spectral quantities. If the underlying processes are stationary, then the correlation functions are diagonal in frequency space. For stationary processes, local error bars can be imposed for spectra in the frequency

domain, whereas the corresponding error bars for correlation functions in the time domain are nonlocal (Percival and Walden, 1993). In addition, if the data contains oscillatory components, which is true for the data treated here, they are compactly represented in frequency space. These reasons form the basis for using a frequency-based representation. For some other advantages of spectra over autocovariance functions, see Percival and Walden (1993), pp. 147–149. The arguments made here are directly applicable to the continuous processes that are of interest in the current paper. Similar arguments also apply to the computation of correlation functions for spike trains. An exception should be made for those spike train examples where there are sharp features in the time domain correlation functions, e.g., due to monosynaptic connections. However, broader features in spike train correlation functions are better studied in the frequency domain.

Despite the advantages of the frequency domain indicated above, the frequent presence of nonstationarity in the data makes it necessary in most cases to use a time-frequency representation. In general, the window for spectral analysis is chosen to be as short as possible to be consistent with the spectral structure of the data, and this window is translated in time. Fundamental to time-frequency representations is the uncertainty principle, which sets the bounds for simultaneous resolution in time and frequency. If the time-frequency plane is ‘tiled’ so as to provide time and frequency resolutions Δt by Δf , then $\Delta t \Delta f \geq 1$. Although there has been a lot of work involving tiling the time-frequency plane using time-scale representations (wavelet bases), we choose to work with frequency rather than scale as the basic quantity, since the time series we are dealing with are better described as having structure in the frequency domain. In particular, the spectra typically have large dynamic range (which indicates good compression of data in the frequency space), and also have spectral peaks, rather than being scale invariant. The time-frequency analysis is often crucial to see this structure, since a long-time average spectrum may prove to be quite featureless. An example of this is presented by MEG data in Section 2.4.1.

1.1.2 Digitization rate, Nyquist frequency, Fourier transforms

Some quantities central to the discussion are defined below. Consider a time series window of length T . The frequency resolution is given by the so-called Raleigh frequency, $\delta f = 1/T$. In all real examples, the time series is obtained at discrete time locations. If we assume the discrete time locations are uniformly spaced at intervals of Δt , then the number of time points N in the interval T is given by $N = T/(\Delta t)$. The digitization frequency or digitization rate is by definition $1/(\Delta t)$. An important quantity is the Nyquist frequency f_{Nyq} , defined as half the digitization frequency $f_{Nyq} = 1/(2\Delta t)$. In this context, it is important to recall the Nyquist theorem and the related concept of aliasing. The basic idea of the Nyquist theorem is that the digitized time series is a faithful representation of the original continuous time series as long as the original one does not contain any frequency components above the Nyquist frequency. Put in a different way, a continuous time series, which is band-limited to the interval $[-B, B]$, should be digitized at a rate $> 2B$ (i.e., $\Delta t = 1/(2B)$) in order to retain all the information present in the continuous time series. Aliasing is an undesirable effect that occurs if this criterion is violated, namely if a time series with frequency content outside the frequency interval $[-f_{Nyq}, f_{Nyq}]$ is digitized at an interval $\Delta t = 2/f_{Nyq}$. The spectral power outside the specified interval is then ‘aliased’ back into the interval $[-f_{Nyq}, f_{Nyq}]$. Consequently, the Nyquist theorem tells us how frequently a continuous time series should be digitized. These concepts are fundamental to the discussion, and the reader unfamiliar with them would benefit from consulting an appropriate text (e.g., (Percival and Walden, 1993)).

The Fourier transform $\tilde{x}(f)$ of a discrete time series $\{x_t | t = n\Delta t, n = 1, 2, \dots, N\}$ is defined in this paper as

$$\tilde{x}(f) = \sum_{n=1}^N x_t \exp(-2\pi i f n \Delta t) \quad (1.1)$$

In places where we use the convention $\Delta t = 1$, the above equation may be rewritten replacing n by t as

$$\tilde{x}(f) = \sum_{t=1}^N x_t \exp(-2\pi i f t) \quad (1.2)$$

The corresponding Nyquist frequency is dimensionless and numerically equal to $1/2$. In any real application, of course, both the digitization rate and the Nyquist frequency have appropriate units. The total time window length T now becomes interchangeable with N . More generally, $T = N\Delta t$ as above. One frequent source of confusion is between the Fourier transform defined above and the fast Fourier transform (FFT). The FFT is an *algorithm* to efficiently compute the Fourier transform on a discrete grid of time points, and should not be confused with the Fourier transform, which is the underlying continuous function of frequency defined above.

1.1.3 Conventional spectral analysis

In this and subsequent subsections, we set $\Delta t = 1$, $T = N$. Physical units are restored where appropriate. We consider below a model example of a time series constructed by adding a stochastic piece, consisting of an autoregressive process excited by white Gaussian noise, to three sinusoids. The time series is given by $y_t = x_t + \sum_{k=1}^3 A_k \sin(2\pi f_k t + \phi_k)$, where $t = 1, 2, \dots, N$ with $N=1024$, and the parameters of the sinusoids are $A_k = (0.7, 0.7, 0.08)$, $f_k = (0.122, 0.391, 0.342)$, and $\phi_k = (0, \pi/3, 2\pi/3)$ for $k = (1, 2, 3)$. Here x_t is an autoregressive process of order 4 given by $x_t = \sum_{k=1}^4 a_k x_{t-k} + \epsilon_t$, with $a_k = (1.87, -1.96, 1.55, -0.683)$. For successive time samples t , ϵ_t are independently drawn from a normal distribution with unit variance. In Fig. 2, the first 300 points of the time series example are plotted.

In conventional nonparametric spectral analysis a tapered Fourier transform of the data is used to estimate the power spectrum. There are various choices of tapers. We use 'taper' rather than 'window', because window is used below to label a segment of data used in spectral analysis, particularly in time-frequency analysis. A taper with optimal band-limiting properties is the zeroth discrete prolate spheroidal sequence, which is described in detail later in the text. Using this taper, a single taper spectral estimate is given in the top right corner of Figure 1.4. In this and the following section, the Fourier transforms are

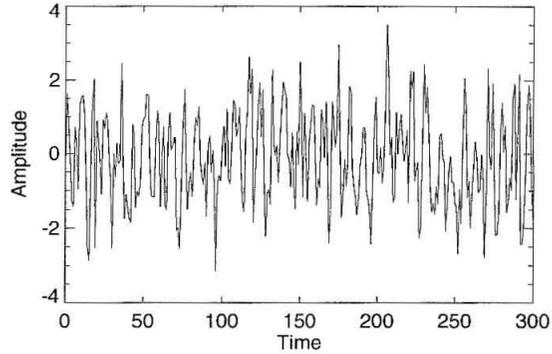


Figure 1.1: Piece of a time series composed of a stochastic piece generated by an order 4 autoregressive process added to a sum of three sinusoids (see text).

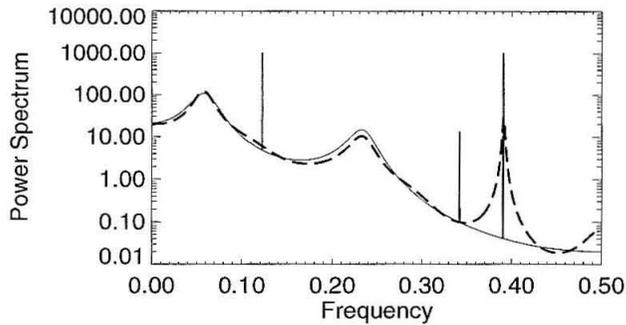


Figure 1.2: Autoregressive spectral estimate of time series data shown in Figure 1.1 (*dashed line*). Also shown is the theoretical spectrum of the underlying process (*solid line*). Note that the sine waves lead to delta functions, whose heights have been determined for the display so that the integrated intensity over a Rayleigh frequency gives the correct strength of the theoretical delta function. The autoregressive estimate completely misses the peaks at $f = 0.122$ and $f = 0.391$.

implemented using an FFT after the time series of length N is padded out to length $4N$ or $8N$. This still gives a discrete representation of the corresponding continuous functions of frequency; however, the grid is sufficiently fine that the resulting function appears to be smooth as a function of frequency.

1.1.4 Autoregressive spectral estimation

To illustrate a parametric spectral estimate, we show in Figure 1.2 the results of an autoregressive (AR) modeling of the data using an order 19 AR process. We used the Levinson-Durbin procedure for purposes of this illustration (Percival and Walden (1993) p.397). The order of the AR model was determined using the Akaike information criterion (Akaike, 1974). Although the parametric estimate is smooth, it fails to accurately estimate the underlying theoretical spectrum. In particular, it completely misses the delta function peaks at $f = 0.122$ and $f = 0.391$.

Some comments are in order regarding AR spectral estimates. The basic weakness of this method is that it starts with the correlation function of the data in order to compute the AR coefficients. This, however, presupposes the answer, since the correlation function is nothing but the Fourier transform of the spectrum: if the correlation function were actually known, there would be no need to estimate the spectrum. In practice, an estimate of the correlation function is made from the data, which contains the same bias problems as in estimating spectra. In fact, in obtaining the illustrated fit, we computed the correlation function by Fourier transforming a direct multitaper spectral estimate (to be described below). Attempts to escape from the circularity pointed out above usually result in strong model assumptions, which then lead to misfits in the spectra (for further discussions, see Thomson (1990), p.614). Despite these problems, AR methods do have some use in spectral estimation, namely to obtain prewhitening filters that reduce the dynamic range of the process, and thus help reduce bias in the final spectral estimate. Another valid usage of AR methods is to treat sufficiently narrow-band signals that can be appropriately modeled by low order AR processes.

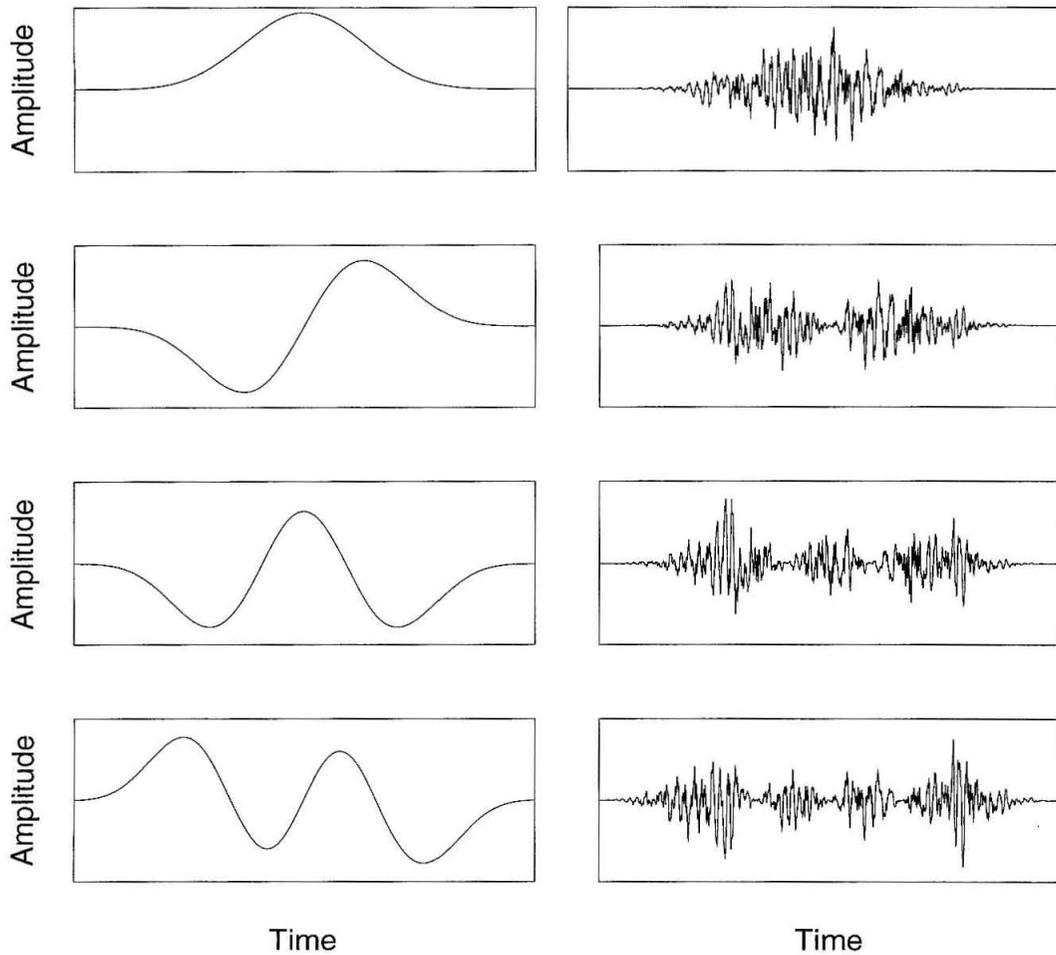


Figure 1.3: *Left column:* prolate spheroidal data tapers for $NW = 5$ ($k=1..4$). *Right column:* tapered time series corresponding to the time series in the example given in the text (part of which is shown in Figure 1.1), multiplied by the data tapers in the left column.

1.1.5 Multitaper spectral analysis

Here we present a brief review of multitaper estimation (Thomson, 1982). This method involves the use of multiple orthogonal data tapers, in particular prolate spheroidal functions, which provide a local eigenbasis in frequency space for finite length data sequences. A summary of the advantages of this technique can be found in Percival and Walden (1993), Chapter 7.

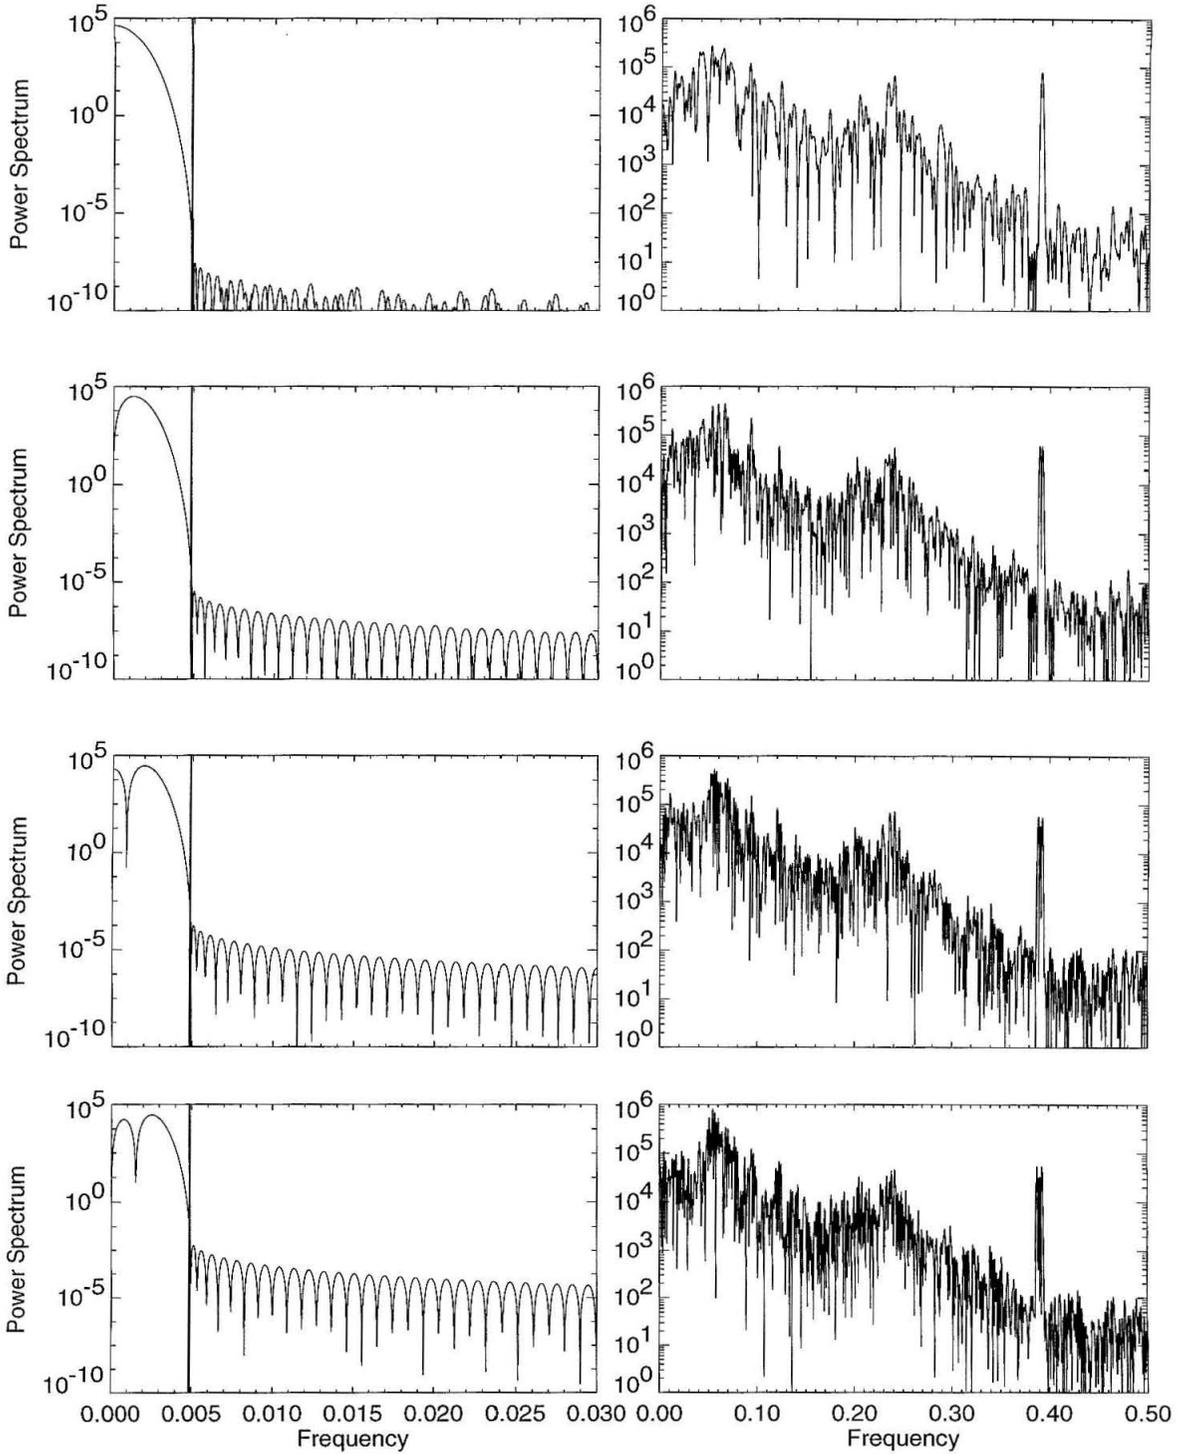


Figure 1.4: *Left column:* spectra of the data tapers shown in Figure 1.3, left column. *Right column:* spectra of the tapered time series shown in Figure 1.3, right column, giving single taper estimates of the spectrum.

Consider a finite length sample of a discrete time process x_t , $t = 1, 2, \dots, N$. Let us assume a spectral representation for the process,

$$x_t = \int_{-\frac{1}{2}}^{\frac{1}{2}} dX(f) \exp(2\pi i f t) \quad (1.3)$$

The Fourier transform of the data sequence $\tilde{x}(f)$ is therefore given by

$$\tilde{x}(f) = \sum_1^N x_t \exp(-2\pi i f t) = \int_{-\frac{1}{2}}^{\frac{1}{2}} K(f - f', N) dX(f') \quad (1.4)$$

where

$$K(f - f', N) = \exp(-2\pi i (f - f')(N + 1)/2) \frac{\sin(N\pi(f - f'))}{\sin(\pi(f - f'))} \quad (1.5)$$

Note that for a stationary process, the spectrum is given by $S(f)df = E[|dX(f)|^2]$. A simple estimate of the spectrum (apart from a normalization constant) is obtained by squaring the Fourier transform of the data sequence, i.e., $|\tilde{x}(f)|^2$. This suffers from two difficulties: first, $\tilde{x}(f)$ is not equal to $X(f)$ except when the data length is infinite, in which case the kernel in Equation 1.5 becomes a delta function. Rather, it is related to $X(f)$ by a convolution, as given by Equation 1.4. This problem is usually referred to as “bias,” and corresponds to a mixing of information from different frequencies of the underlying process due to a finite data window length. Second, even if the data window length were infinite, calculating $|\tilde{x}(f)|^2$ without using a tapering function (a quantity known as the periodogram) effectively squares the observations without averaging; the spectrum is the *expectation* of this squared quantity. This issue is referred to as the lack of consistency of the periodogram estimate, namely the failure of the periodogram to converge to the spectrum for large data lengths. The reason for this is straightforward. When one takes a fast Fourier transform of the data, one is estimating N quantities from N data values, which obviously leads to overfitting if the data are stochastic. More precisely, the squared Fourier transform of the

time series is an inconsistent estimator of the spectrum, because it does not converge as the data time series tends to infinite length.

To resolve the first issue, the data are usually multiplied by a data taper, which leads to replacing the kernel in Equation 1.5 by a kernel that is more localized in frequency. However, this leads to the loss of the ends of the data. To surmount the second problem, the usual approach is to average together overlapping segments of the time series (Welch, 1967). Repetition of the experiment also gives rise to an ensemble over which the expectation can be taken, but we are interested in single trial calculations involving only a single time series. Evidently, some amount of smoothing is necessary to reduce the variance of the estimate, the question being what is an appropriate and systematic way to performing this smoothing.

An elegant approach toward the solution to both of the above problems is provided by the multitaper spectral estimation method in which the data are multiplied by not one, but several orthogonal tapers, and Fourier transformed to obtain the basic quantity for further spectral analysis. The method can be motivated by treating Equation 1.4 as an integral equation to be solved in a regularized way. The simplest example of the method is given by the direct multitaper spectral estimate $S_{MT}(f)$, defined as the average over individual tapered spectral estimates,

$$S_{MT}(f) = \frac{1}{K} \sum_{k=1}^K |\tilde{x}_k(f)|^2 \quad (1.6)$$

where

$$\tilde{x}_k(f) = \sum_1^N w_t(k) x_t \exp(-2\pi i f t) \quad (1.7)$$

Here $w_t(k)$ ($k = 1, 2, \dots, K$) are K orthogonal taper functions with appropriate properties. A particular choice of these taper functions, with optimal spectral concentration properties, is given by the discrete prolate spheroidal sequences (DPSS) (Slepian and Pollak, 1961). Let $w_t(k, W, N)$ be the k th DPSS of length N and frequency bandwidth pa-

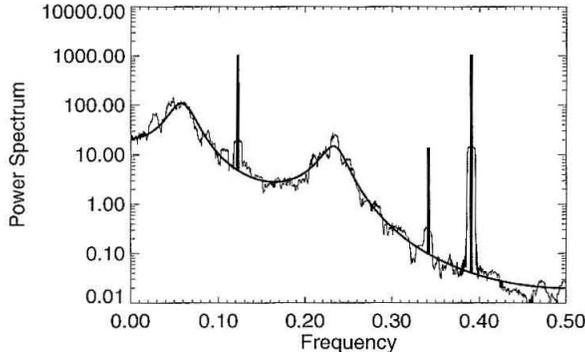


Figure 1.5: Direct multitaper spectral estimate of time series example using $WT = 5$, $K = 9$. The estimates for the first four of the nine tapers that are averaged to create this estimate are shown in Figure 1.4. The theoretical spectrum is also shown.

parameter W . The DPSS form an orthogonal basis set for sequences of length, N , and are characterized by a bandwidth parameter W . The important feature of these sequences is that for a given bandwidth parameter W and taper length N , $K = 2NW - 1$ sequences out of a total of N each have their energy effectively concentrated within a range $[-W, W]$ of frequency space. Consider a sequence w_t of length N whose Fourier transform is given by $U(f) = \sum_1^N w_t \exp(-2\pi i f t)$. Then we can consider the problem of finding sequences w_t so that the spectral amplitude $U(f)$ is maximally concentrated in the interval $[-W, W]$, i.e.,

$$\int_{-W}^W |U(f)|^2 df \quad (1.8)$$

is maximized, subject to a normalization constraint which may be imposed using a Lagrange multiplier. It can be shown that the optimality condition leads to a matrix eigenvalue equation for $w_t(k, W, N)$

$$\sum_{t'=1}^N \frac{\sin[2\pi W(t-t')]}{\pi(t-t')} w_{t'} = \lambda w_t \quad (1.9)$$

The eigenvectors of this equation are the DPSS. The remarkable fact is that the first $2NW$ eigenvalues $\lambda_k(N, W)$ (sorted in descending order) are each approximately equal to

one, while the remainder are approximately zero. Since it follows from the above definitions that

$$\lambda_k(N, W) = \frac{\int_{-W}^W |U(f)|^2 df}{\int_{-\frac{1}{2}}^{\frac{1}{2}} |U(f)|^2 df}, \quad (1.10)$$

this is a precise statement of the spectral concentration mentioned above. The fact that many of the eigenvalues are close to one makes the eigenvalue problem Equation 1.9 ill-conditioned and unsuitable for the actual computation of the prolates [this can be achieved by a better-conditioned tridiagonal form (Percival and Walden, 1993)]. The DPSS can be shifted in concentration from $[-W, W]$ centered around zero frequency to any non-zero center frequency interval $[f_0 - W, f_0 + W]$ by simply multiplying by the appropriate phase factor $\exp(2\pi i f_0 t)$, an operation known as demodulation. The usual strategy is to select the desired analysis half-bandwidth W to be a small multiple of the Raleigh frequency $1/N$, and then take the leading $2NW - 1$ DPSS as data tapers in the multitaper analysis. Note that $\leq 2NW$ of the sequences are typically taken, since the last few of these have progressively worsening spectral concentration properties.

For illustration, in the left column of Figure 1.3 we show the first 4 DPSS for $W = 5/N$. Note that the orthogonality condition ensures that successive DPSS each have one more zero crossing than the previous one. In the right column of Figure 1.3, we show the time series example from the earlier subsection multiplied by each of the successive data tapers. In the left column of Figure 1.4 we show the spectra of the data tapers themselves, showing the spectral concentration property. The vertical marker denotes the bandwidth parameter W .

In Figure 1.4 we show the magnitude squared Fourier transforms of the tapered time series shown in Figure 1.3. The arithmetic average of these spectra for $k = 1, 2, \dots, 9$ (note that only 4 of 9 are shown in Figures 1.3 and 1.4) gives a direct multitaper estimate of the underlying process, shown in Figure 1.5. Also shown in that figure is the theoretical spectrum of the underlying model.

In the direct estimate, an arithmetic average of the different spectra is taken. However,

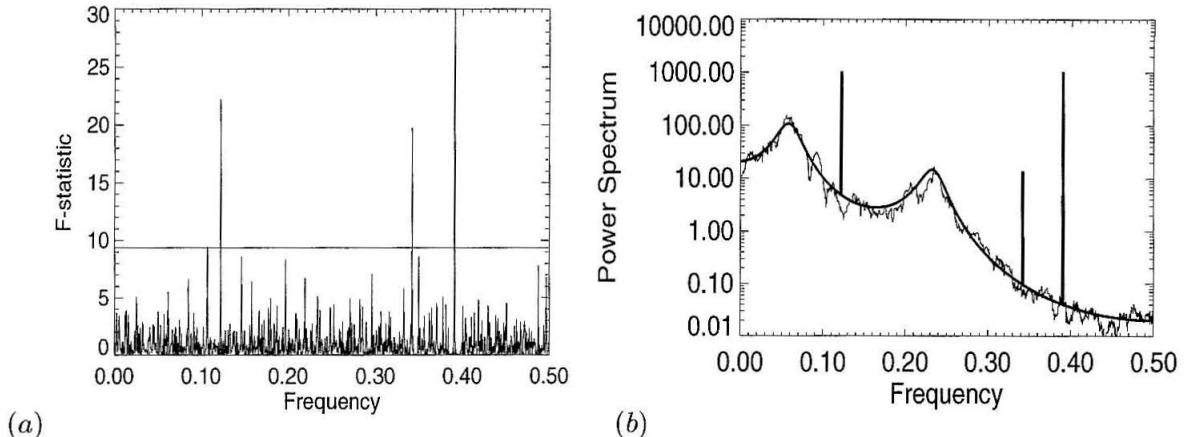


Figure 1.6: *a)* F-statistic values for testing for the presence of a significant sinusoidal component in the time series example along with significance levels. *b)* Direct multitaper spectral estimate of the time series example using $WT = 5$, $K = 9$ reshaped to remove the sinusoidal lines. The theoretical spectrum is also shown in bold.

the different data tapers differ in their spectral side-lobes, so that a weighted average is more appropriate. In addition, the weighting factors should be chosen adaptively depending on the local variations in the spectrum. For more detailed considerations along these lines, the reader is referred to Thomson (1982).

Sine waves in the original time series correspond to square peaks in the multitaper spectral estimate. This is usually an indication that the time series contains a sinusoidal component along with a broad background. The presence of such a sinusoidal component may be detected by a test based on a goodness-of-fit F-statistic (Thomson, 1982). To proceed, let us assume that the data contain a sinusoid of complex amplitude μ at frequency f_0 (the corresponding real series being $\text{Re}(\mu \exp(2\pi i f_0 t))$). Let us also assume that in a frequency interval $[f_0 - W, f_0 + W]$ around f_0 , the process, to which the sinusoid is added, is white. Note that this is a much weaker assumption than demanding that the process is white over the entire frequency range. Under these assumptions, for $f \in [f_0 - W, f_0 + W]$, the tapered Fourier transforms of the data are given by

$$\tilde{x}_k(f) = \mu U_k(f - f_0) + n_k(f), \quad k = 1, 2, \dots, K \quad (1.11)$$

Here $U_k(f)$ is the Fourier transform of the k th DPSS, and f is in the range $[f_0 - W, f_0 + W]$. The assumption of a locally white background implies that $n_k(f)$ are independently and identically distributed complex Gaussian variables. Treating Equation 1.11 as a linear regression equation at $f = f_0$ leads to an estimate of the sine wave amplitude (which corresponds to a particular tapered Fourier transform of the data)

$$\hat{\mu}(f_0) = \frac{\sum_k \tilde{x}_k(f_0) U_k^*(0)}{\sum_k |U_k(0)|^2} \quad (1.12)$$

and to an F statistic for the significance of a non-zero μ

$$F(f_0) = \frac{|\hat{\mu}(f_0)|^2}{\sum_k |\tilde{x}_k(f_0) - \hat{\mu}(f_0) U_k(0)|^2} \quad (1.13)$$

Under the null hypothesis that there is no line present, $F(f_0)$ has an F distribution with $(2, 2K - 2)$ degrees of freedom. We plot this function $F(f)$ for the time series in the above example in Figure 1.6a. For this example we have chosen $W = 7/N$, $K = 13$. One obtains an independent F statistic every Raleigh frequency, and since there are N Raleigh frequencies in the spectrum, the statistical significance level is chosen to be $1 - 1/N$. This means that on an average, there will be at most one false detection of a sinusoid across all frequencies. A horizontal line in Figure 1.6a indicates this significance level. Thus, a line crossing this level is found to be a significant sinusoid present in the data. The sinusoids known to be present in the data are shown to give rise to very significant F-statistics at this level of significance, and there is one spurious crossing of the threshold by a small amount. The linear regression leads to an estimate for the sinusoid amplitudes. In the present example, the percentage differences between the original and estimated amplitudes were found to be 6%, 4%, and 2% for increasing frequencies; the corresponding errors for the phase were 0.2%, 2%, and 2%. Note that the errors in phase estimation are smaller than the errors in the amplitude estimates. From the estimated amplitudes and phases, the

sinusoidal components can be reconstructed and subtracted from the data. The spectrum of the residual time series can then be estimated by the same techniques. This residual spectrum is shown in Figure 1.6b, along with the theoretical spectrum of the underlying autoregressive process.

1.1.6 Choice of the bandwidth parameter

The choice of the time window length $T = N\Delta t$ and the bandwidth parameter W is critical for applications. No simple procedure can be given for these choices, because the choice really depends on the data set at hand, and is best made iteratively by visual inspection and some degree of trial and error. $2TW = 2W/(1/T)$ gives the number of Raleigh frequencies over which the spectral estimate is effectively smoothed, so that the variance in the estimate is typically reduced by $2TW$. Thus, the choice of W is a choice of how much to smooth. In qualitative terms, the bandwidth parameter should be chosen to reduce variance while not overly distorting the spectrum. This can be done formally by trading off an appropriate weighted sum of the estimated variance and bias. However, as a rule of the thumb we find fixing the time bandwidth product TW at a small number (typically 3 or 4) and then varying the window length in time until sufficient spectral resolution is obtained is a reasonable strategy. This presupposes that the data are examined in the time frequency plane so that T may be significantly smaller than the total data length.

1.2 Analysis of multivariate data

So far we have concentrated on the analysis of univariate time series. However, the principal subject of this paper is multichannel data, so we now consider the analysis of multivariate time series. The basic methods for dealing with such data are similar to those used for other multivariate data, with modifications to take into account the fact that we are dealing with time series. In fact, a scalar time series can itself be usefully represented as a multivariate time series by going to a lag-vector representation or a time-frequency representation. Such a representation may be desirable to understand the structure underlying the scalar time

series. Other examples of multichannel data include multiple spike trains and various forms of brain imaging, including optical imaging using intrinsic and extrinsic contrast agents, magnetic resonance imaging, and magnetoencephalography. In general, one can think of one or two space dimensions added to the time dimension in the data. In the sections below, we briefly review some of the concepts useful to our analysis later in the paper. The techniques can be grouped into two general classes: choosing an appropriate low-dimensional representation (e.g., in principal components analysis) or choosing a partitioning of the multidimensional space (e.g., in various forms of clustering). In this section we concentrate on mode decomposition. A discussion of clustering methods may be found in Duda and Hart (1973).

1.2.1 Eigenmode analysis: SVD

The singular value decomposition (SVD) is a representation of a general matrix of fundamental importance in linear algebra that is widely used to generate canonical representations of multivariate data. It is equivalent to principal component analysis in multivariate statistics, but in addition is used to generate low-dimensional representations for complex multidimensional time series. The SVD of an arbitrary (in general complex) $p \times q$ matrix ($p > q$) \mathcal{M} is given by $\mathcal{M} = U\Lambda V^\dagger$, where the $p \times q$ matrix U has orthonormal columns, the $q \times q$ matrix Λ is diagonal with real, non-negative entries and the $q \times q$ matrix V is unitary. Note that the matrices $\mathcal{M}\mathcal{M}^\dagger = U\Lambda^2U^\dagger$ and $\mathcal{M}^\dagger\mathcal{M} = V\Lambda^2V^\dagger$ are hermitian, with eigenvalues corresponding to the diagonal entries of Λ^2 and U and V the corresponding matrices of eigenvectors. Consider the special case of space-time data $I(\mathbf{x}, t)$. The SVD of such data is given by

$$I(\mathbf{x}, t) = \sum_n \lambda_n I_n(\mathbf{x}) a_n(t) \quad (1.14)$$

where $I_n(\mathbf{x})$ are the eigenmodes of the "spatial correlation matrix"

$$C(\mathbf{x}, \mathbf{x}') = \sum_t I(\mathbf{x}, t) I(\mathbf{x}', t) \quad (1.15)$$

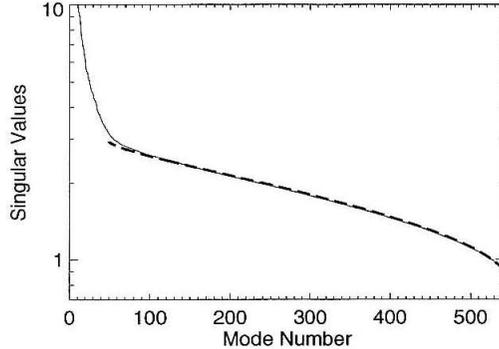


Figure 1.7: Sorted singular values corresponding to a space-time SVD of functional MRI image time series from data set \mathcal{D} . The tail in the spectrum is fit with the theoretical formula given in the text.

Similarly, $a_n(t)$ are the eigenmodes of the "temporal correlation matrix"

$$C(t, t') = \sum_{\mathbf{x}} I(\mathbf{x}, t) I(\mathbf{x}, t') \quad (1.16)$$

If the sequence of images were randomly chosen from an ensemble of spatial images, then $C(\mathbf{x}, \mathbf{x}')$ would converge to the ensemble spatial correlation function in the limit of many time samples. If in addition the ensemble had space translational invariance, then the eigenmodes $I_n(\mathbf{x})$ would be plane waves $\exp(i\mathbf{k} \cdot \mathbf{x})$, the mode number "n" would correspond to wave vectors and the singular values would correspond to the spatial structure factor $S(\mathbf{k})$. In general, the image ensemble in question will not have translational invariance; however, the SVD will then provide a basis set analogous to wave vectors. In physics one normally encounters structure factors $S(\mathbf{k})$ that decay with wave vector. In the more general case, the singular value spectrum, organized in descending order, will show a decay indicative of the structure in the data.

To make sense of an SVD performed on a data matrix, it is important to know the expected distribution of singular values if the entries of the matrix were random. This problem lies in the domain of random matrix theory, and can be solved in special cases (Sengupta and Mitra, 1999). As an example, consider the case of a $p \times q$ matrix $\mathcal{M} = \mathcal{M}_0 + \mathcal{N}$

where \mathcal{M}_0 is fixed and the entries of N are independently normally distributed with zero mean and identical variance σ^2 . \mathcal{M}_0 may be thought of as the desired or underlying signal; for an SVD to be useful, \mathcal{M}_0 should effectively have a low rank structure. A typical procedure is to take the SVD of \mathcal{M} and to truncate the singular value spectrum to keep only values that cross a threshold. Consider the special case in which $\mathcal{M}_0 = 0$, so that we are dealing with a purely noise matrix. In this case, it can be shown that (Denby and Mallows, 1991; Sengupta and Mitra, 1999) the density of singular values, defined as

$$\rho(\lambda) = \left\langle \sum_n \delta(\lambda - \lambda_n) \right\rangle \quad (1.17)$$

is given in the limit of large matrix sizes by

$$\rho(\lambda) = \frac{1}{\pi\sigma\lambda} \sqrt{(\lambda_+^2 - \lambda^2)(\lambda^2 - \lambda_-^2)} \quad (1.18)$$

where (Recall that σ^2 is the variance of the matrix entries, and p, q are the dimensions of the matrix)

$$\lambda_{\pm}^2 = 2\sigma^2 \left[\frac{p+q}{2} \pm \sqrt{pq} \right] \quad (1.19)$$

It is somewhat easier to work with the integrated density of states $P(\lambda) = \int_0^\lambda \rho(x)dx$, since λ plotted against $1 - P(\lambda)$ gives the sorted singular values (in decreasing order). More generally, \mathcal{M}_0 is not zero but is given by a low rank matrix. The distribution of the singular values can be worked out in this case, but if the original matrix has low rank and if the ‘signal’ singular values are large compared to the ‘noise’ singular values, then the singular value distribution shows a tail which can be fit by the above formula, the only quantity needing adjustment being the total weight under the density function $\rho(\lambda)$ (i.e., an overall normalization factor).

To illustrate the above, in Figure 1.7 we show the sorted singular values from the SVD of data set \mathcal{C} consisting of 550 frames of fMRI data. The original image data consisted of 550 images each 64×64 pixels big, but the space data was first masked to select out

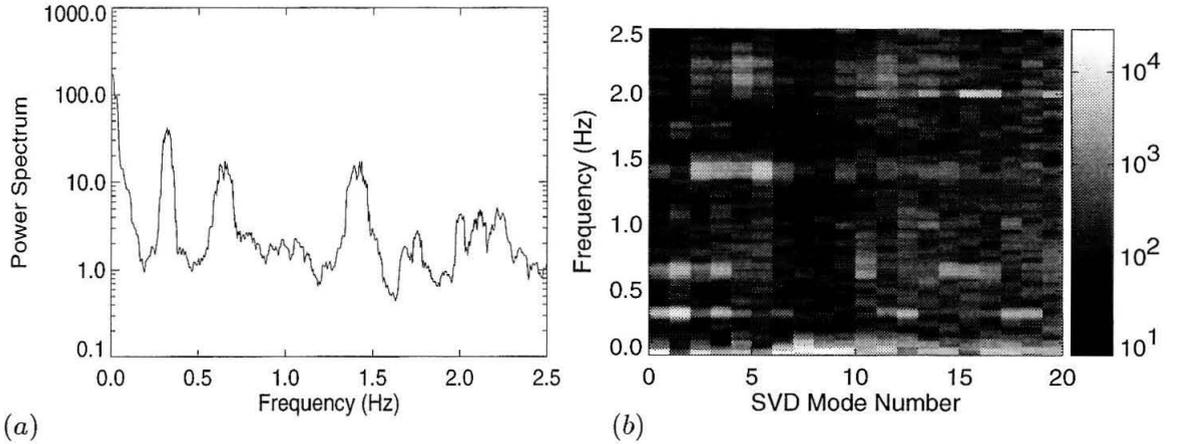


Figure 1.8: Spectra of first 20 principal component time series from data set \mathcal{D} . (a) Weighted average spectrum. (b) Image showing the spectra versus mode number. The log amplitudes of the spectra are color coded.

1877 pixels. Thus, the SVD is performed on a 1877×550 matrix. The resulting singular values are shown with the range truncated to magnify the noise tail. Also shown is the theoretical singular value spectrum expected for the noise tail (dashed line) based on the formula Equation 1.18 for a pure noise matrix with a single adjustable parameter σ^2 , which has been selected to match the middle portion of the tail. The total weight of the density has been adjusted to account for the last 500 singular values.

Unlike in the temporal domain, where going to a frequency-based representation does make sense for neurobiological data, the spatial wave-vector representation is not of general use because of the generic lack of translational invariance in space. However, the spatial basis generated by an SVD is somewhat more meaningful. It may, for example, reflect underlying anatomical structure. Application of the SVD on space-time imaging data may be found in the literature, sometimes with modifications. However, the space-time SVD suffers from a severe drawback in the present context. The difficulty is that there is no reason why the neurobiologically distinct modes in the data should be orthogonal to each

other, a constraint imposed by the SVD. In practice, it is observed that an SVD on space-time data, different sources of fluctuations, such as cardiac and respiratory oscillations and results of neuronal activity, may appear in the same mode of the decomposition, thus preventing an effective segregation of the different degrees of freedom.

As an example, consider the SVD of the fMRI data set \mathcal{C} , the corresponding singular value distribution of which has been shown in Figure 1.7. Note that in these data the digitization rate was 5 Hz, and the length of the time series 110 s. The mixing of physiologically distinct processes in the individual principal component (PC) time series thus obtained can be seen by studying the spectra of the PCs. Figure 1.8a shows the average spectrum across PCs of data set \mathcal{C} . The peak near zero frequency corresponds to the stimulus response as well as possible vasomotor oscillation or other slow fluctuations. The peaks at 0.3 Hz and 0.6 Hz correspond to breathing, and the peak at 1.3 Hz corresponds to the cardiac cycle. In Figure 1.8b, the spectra are shown for the first 20 modes (with the largest singular values). The spectra are coded as gray scale intensities, and are shown against the corresponding mode numbers. As is clear from studying the spectra as a function of mode number (Fig. 1.8b), the decomposition mixes the various effects.

We describe below a more effective way of separating distinct components in the image time series dealt with here using a decomposition analogous to the space-time SVD, but in the space-frequency domain. The success of the method stems from the fact that the data in question are better characterized by a frequency based representation.

1.2.2 Space-frequency SVD

From the presence of spectral peaks in Figure 1.8a it can be inferred that different components in the dynamic image data may be separated if the SVD were performed after localizing the data in the frequency domain. This can be achieved by projecting the space-time data to a frequency interval, and then performing SVD on this space-frequency data (Thomson and Chave, 1991; Mann and Park, 1994; Mitra et al., 1997). Projecting the data on a frequency interval can be performed effectively by using DPSS with the appropriate bandwidth parameter. For a fixed center frequency f and a half bandwidth W , consider

the projection matrix

$$P_{k,t}(f; W) = w_k(t, W) \exp(2\pi itf) \quad k = 1, 2, \dots, K \quad (1.20)$$

In the above, we assume that $K = [2NW]$, where $[X]$ is the integer closest to X but less than X . Acting on the space of sequences of length N , this matrix projects out a subspace with frequencies concentrated in $[f - W, f + W]$. Note that $P^\dagger P$ serves as an optimal bandlimiting filter on the time series. Given the $N_x \times N$ space-time data matrix $I = I(x, t)$, the space-frequency data corresponding to the frequency band $[f - W, f + W]$ are given by the $N_x \times K$ complex matrix $\tilde{I} = IP^T$. In expanded form,

$$\tilde{I}(x, k; f) = \sum_{t=1}^N I(x, t) w_k(t, W) \exp(2\pi itf) \quad (1.21)$$

We are considering here the SVD of the $N_x \times K$ complex matrix with entries $\tilde{I}(x, k; f)$ for fixed f .

$$\tilde{I}(x, k; f) = \sum_n \lambda_n(f) \tilde{I}_n(x; f) a_n(k; f) \quad (1.22)$$

This SVD can be carried out as a function of the center frequency f , using an appropriate choice of W . In the best case most of the coherent structure is captured by the dominant singular vector at each frequency. At each frequency f one obtains a singular value spectrum $\lambda_n(f)$ ($n=1,2,\dots,K$), the corresponding (in general complex) spatial mode $\tilde{I}_n(x; f)$, and the corresponding local frequency modes $\tilde{a}_n(k; f)$. The frequency modes can be projected back into the time domain using $P_{k,t}(f; W)$ to give (narrow-band) time varying amplitudes of the complex eigenimage.

In the space-frequency SVD computation, an overall coherence $\bar{C}(f)$ may be defined as (it is assumed that $K \leq N_x$)

$$\bar{C}(f) = \frac{\lambda_1^2(f)}{\sum_{n=1}^K \lambda_n^2(f)} \quad (1.23)$$

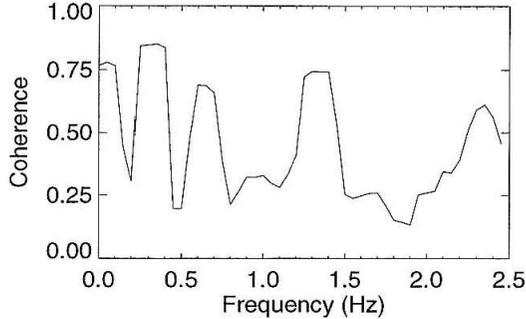


Figure 1.9: Overall coherence spectrum corresponding to the functional MRI time series examined in Figures 1.7 and 1.8.

The overall coherence spectrum then reflects how much of the fluctuation in the frequency band $[f - W, f + W]$ is captured by the dominant spatial mode. If the image data are completely coherent in that frequency band, then $\bar{C}(f) = 1$. More generally, $1 \geq \bar{C}(f) \geq 0$, and for random data, assuming $N_x \gg K$, $\bar{C}(f) \sim \frac{1}{K}$. If N_x and K are comparable, then results such as those in the previous section may be used to determine the distribution of $\bar{C}(f)$.

To illustrate this technique, we show results of its application to data set \mathcal{C} of fMRI data. The calculation used 19 DPSS, corresponding to a full bandwidth of 0.2 Hz. The overall coherence spectrum resulting from a space-frequency SVD analysis of this data is shown in Figure 1.9. Note the correspondence of this spectrum, which is dimensionless, to the power spectrum presented in Figure 1.8a. The magnitudes $|\tilde{I}_1(x; f)|$ of the dominant spatial eigenmodes as a function of frequency are shown in Figure 1.10. The leading eigenmodes separate the distinct sources of fluctuations as a function of frequency.

1.3 Local frequency ensemble and jackknife error bars

One important advantage of the multitaper method is that it offers a natural way of estimating error bars corresponding to most quantities obtained in time series analysis, even if one is dealing with an individual instance of a time series. The fundamental notion here is that of

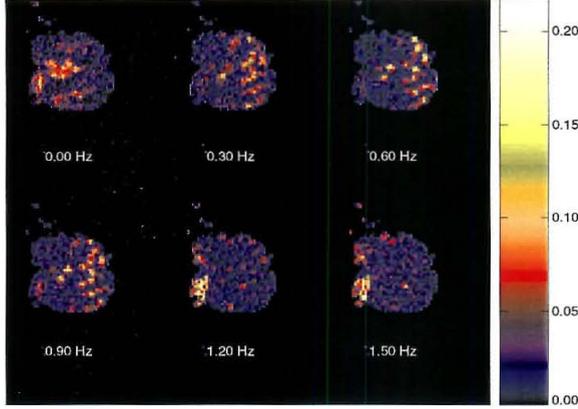


Figure 1.10: Amplitudes of leading spatial eigenmodes for the space-frequency SVD of fMRI data from data set \mathcal{C} . Note that the spatial structure varies as a function of center frequency with the physiological oscillations segregating into distinct frequency bands.

a *local frequency ensemble*. The idea is that if the spectrum of the process is locally flat over a bandwidth $2W$, then the tapered Fourier transforms $\tilde{x}_k(f) = \sum_{t=1}^N \exp(-2\pi i f t) x(t) w_k(t)$ constitute a statistical ensemble for the Fourier transform of the process at the frequency f_0 . Assuming that the underlying process is locally white in the frequency range $[f_0 - W, f_0 + W]$, then it follows from the orthogonality of the data tapers that $\tilde{x}_k(f)$ are uncorrelated random variables with the same variance. For large N , $\tilde{x}_k(f)$ may be assumed to be asymptotically normally distributed under some general circumstances (for related results see Mallows (1967)). This provides one way of thinking about the direct multitaper estimate presented in the previous sections: the estimate consists of an average over the local frequency ensemble.

The above discussion serves as a motivation for multitaper estimates of the correlation function, the transfer function, and the coherence between two time series. Given two time series x_t, y_t , and the corresponding multiple tapered Fourier transforms $\tilde{x}_k(f), \tilde{y}_k(f)$, the following direct estimates can be defined (Thomson, 1982) for the correlation function $C_{yx}(f) = E[\tilde{y}(f)\tilde{x}^*(f)]$, the transfer function $T_{yx}(f) = E[\tilde{y}(f)\tilde{x}^*(f)]/E[|\tilde{x}(f)|^2]$, and the coherence function $\rho_{yx}(f) = E[\tilde{y}(f)\tilde{x}^*(f)]/\sqrt{E[|\tilde{x}(f)|^2]E[|\tilde{y}(f)|^2]}$:

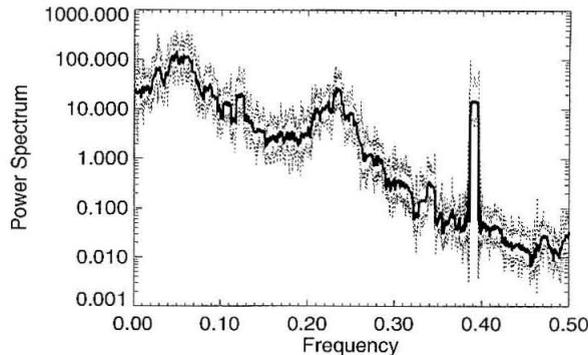


Figure 1.11: Jackknife error bars on multitaper spectral estimate corresponding to spectrum shown in Figure 1.5. The spectral estimate is the solid bar and the error bars are represented by dots.

$$\hat{C}_{yx}(f) = \frac{1}{K} \sum_{k=1}^K \tilde{y}_k(f) \tilde{x}_k^*(f) \quad (1.24)$$

$$\hat{T}_{yx}(f) = \frac{\sum_{k=1}^K \tilde{y}_k(f) \tilde{x}_k^*(f)}{\sum_{k=1}^K \tilde{x}_k(f) \tilde{x}_k^*(f)} \quad (1.25)$$

$$\hat{\rho}_{yx}(f) = \frac{\sum_{k=1}^K \tilde{y}_k(f) \tilde{x}_k^*(f)}{\sqrt{\sum_{k=1}^K |\tilde{x}_k(f)|^2 \sum_{k=1}^K |\tilde{y}_k(f)|^2}} \quad (1.26)$$

These definitions allow the estimation of the coherence and transfer function from a single instance of a pair of time series. By using the local frequency ensemble one can also estimate jackknife error bars for the spectra and the above quantities. The idea of the jackknife is to create different estimates by leaving out a data taper in turn. This creates a set of estimates from which an error bar may be computed (Thomson and Chave, 1991). As an example, we show in Figure 1.11 jackknife estimates of the standard deviations of the spectral estimate in Figure 1.5. Jackknife error bars and transfer functions for spectral estimates averaged across trials can be generated in the same way from the joint ensemble indexed by trial and taper. This procedure is used in the analysis of data in Chapters 3 and 4.

1.4 Spectral analysis of point processes

We now consider point processes and the mathematical description of the firing of action potentials by neurons. We are interested in characterizing second order correlations using spectral analysis and the procedure is similar in spirit to the analysis of continuous time series described above. Point processes can be represented by the number of events in time (the counting process) or by the sequence of time intervals between events (the interval process). These representations are equivalent but complementary (Daley and Vere-Jones, 1988). We give the definitions and formulae for second moments in each representation and then discuss the Poisson process and the renewal process which are the fundamental counting and interval processes respectively. Results exist for univariate and multivariate processes, but here we will exclusively present the univariate case. Univariate results for counting processes can be generalized to the multivariate case without difficulty, but those for interval processes cannot be generalized due to the problem of specifying relationships between intervals in multidimensional time. Our approach follows that of Cox and Lewis (1966).

1.4.1 Counting process

The counting process is given in terms of the number of events, $dN(t)$, that occur in an interval $(t, t + \Delta t)$. Multiple events are excluded by taking the limit $\Delta t \rightarrow 0$ so $dN(t)$ takes the values 0 and 1. The following presents steps in the calculation that lead to the second-order properties of the counting process. The mean and variance of $dN(t)$ are

$$\begin{aligned}
 E[dN(t)] &= \text{prob}\{dN(t) = 1\} \\
 &= \mu_n \Delta t \\
 \text{Var}[dN(t)] &= E[\{dN(t)\}^2] - \{E[dN(t)]\}^2 \\
 &= \text{prob}\{dN(t) = 1\} - [\text{prob}\{dN(t) = 1\}]^2 \\
 &= \mu_n \Delta t
 \end{aligned}$$

where μ_n is the average number of events in the interval $(0, t]$.

Covariance properties of the counting process result from considering

$$\begin{aligned} E[dN(t)dN(t + \tau)] &= \text{prob}\{dN(t) = dN(t + \tau) = 1\} \\ &= \mu_n(\Delta t)\text{prob}\{dN(t + \tau) = 1 | dN(t) = 1\} \end{aligned}$$

The probability on the right can be expressed as the renewal density, $\eta_n(\tau)$, which gives the probability of an event at time $t + \tau$ in the future given an event at time t . This then leads to the following expression for the covariance

$$\text{Cov}[dN(t), dN(t + \tau)] = \mu_n(\Delta t)^2 \{ \eta_n(\tau) - \mu_n \}$$

The covariance function, $C_{dN}(\tau)$, also known as the auto-correlation function is then

$$C_{dN}(\tau) = \mu_n \{ \eta_n(\tau) - \mu_n \} \tag{1.27}$$

The spectrum of a point process, $S_{dN}(f)$, is the Fourier transform of the covariance, or auto-correlation, function:

$$S_{dN}(f) = 1/2\pi \int_{-\infty}^{\infty} C_{dN}(\tau) \exp(-2\pi i f \tau) d\tau \tag{1.28}$$

As is the case for continuous time series, the spectrum and auto-correlation function of a point process contain equivalent information, but the spectrum has better statistical properties. Multitaper methods can be used to estimate the spectrum of a point process.

1.4.2 Interval process

The interval process is defined by the sequence X_1, X_2, \dots of intervals between successive events. This sequence is stationary in time, where time is given by the serial number of the event. This representation is treated in the same way as a continuous time series. The mean and variance of X_i are denoted by μ_x and σ_x^2 , respectively. The ratio, μ_x^2/σ_x^2 , is called the coefficient of variation and measures the deviation of the distribution from an exponential. The correlation properties of the sequence are given by the autocorrelation function, $C_X(k)$, and the spectrum, $S_X(q)$, known here as the interval spectrum to distinguish it from the spectrum of the counts presented above.

The autocorrelation function is given by

$$C_X(q) = \text{cov}(X_i, X_{i+q})/\sigma_x^2 \quad (1.29)$$

and the interval spectrum is given by

$$S_X(k) = \sigma_x^2/(2\pi) \sum_{q=-\infty}^{\infty} e^{-2\pi i k q} C_X(q) \quad (1.30)$$

Multitaper methods can be used to estimate the interval spectrum.

1.4.3 Poisson process

The simplest counting process is the Poisson process which is often used to model random spiking activity. This process is parameterized by a rate: the average number of events occurring during a given interval (Cox and Lewis, 1966). Since the probability of an event occurring at a given time is not dependent on activity before that time, there is no temporal structure in a Poisson process. The spectrum measures dependencies in the temporal ordering of events in a random process. A Poisson process has a flat or white spectrum so spike spectra with significant deviations from uniformity are evidence of temporal structure. If suppression is also present in the spectrum, a Poisson process with a stochastic, time-varying rate cannot describe the activity either (Brillinger, 1978). We explore this issue in

more detail in Chapter 4.

1.4.4 Renewal process

Spike activity that is non-Poisson is often modeled using a renewal process. A renewal process is completely described by its renewal density or ISI histogram; therefore, the temporal structure in a renewal process is determined by the rate of each interval and does not extend beyond neighboring pairs of events. In Chapter 4 we use the interval spectrum, which is sensitive to the ordering of intervals in time, to test whether spiking activity is a renewal process.

1.5 Spectral analysis of hybrid processes

The data considered in Chapter 3 are spiking activity (a point process) and the LFP (a continuous valued time series). The analysis of such 'hybrid' data presents some novel features that we discuss here. Similarly to the situation for point and continuous processes, spectral analysis provides a unified framework for the characterization of correlations in hybrid processes and we use multitaper methods of spectral estimation developed by Thomson (1982) to construct estimators for all spectral quantities. Correlation function measures such as the auto- and cross-correlation function characterize the same statistical structure in time series as spectra and cross-spectra; however, spectral estimates offer significant advantages over their time domain counterparts. An important one is that nearby points in time are highly correlated while neighboring points in frequency are almost independent. The use of data tapers makes estimates of the spectrum even less susceptible than the auto-correlation function to such problems of statistical bias and variance (Brillinger, 1978) and therefore more suitable for investigating temporal structure in experimental data.

The coherency, $C(f)$, is a complex quantity that is a function of frequency for a given window in time and measures the degree of predictability of one process using a linear function of the other (Brillinger, 1975; Rosenberg et al., 1989), or phase coherence. In Chapter 3 we use a hybrid version of the coherency for point process and ordinary time

series data to measure the relation between spiking and LFP activity.

The magnitude of the coherency lies between zero and one and is called the *coherence*. Coherence indicates the strength of a linear relationship between different processes. The coherency is used in preference to the cross-correlation function to measure the relation between two processes. The cross-correlation function is typically normalized by the zero lag estimates of the autocorrelation function. This is *ad-hoc* since it depends on the bin size used to estimate the zero lag autocorrelation. In contrast, the coherency is normalized independently for each frequency band by dividing by the spectrum of each process. This normalization means changes in coherence are not a result of changes in firing rate and allows the coherency to be meaningfully averaged across different pairs of time series.

Chapter 2 Analysis of dynamic brain imaging data

2.1 Introduction

Modern imaging techniques for probing brain function, including functional magnetic resonance imaging, intrinsic and extrinsic contrast optical imaging, and magnetoencephalography, generate large data sets with complex content. This chapter develops appropriate techniques for analysis and visualization of such imaging data, to separate the signal from the noise and characterize the signal. The techniques developed fall into the general category of multivariate time series analysis, and in particular we extensively use the multitaper framework of spectral analysis. We develop specific protocols for the analysis of fMRI, optical imaging and MEG data, and illustrate the techniques by applications to real data sets generated by these imaging modalities. In general, the analysis protocols involve two distinct stages: 'noise' characterization and suppression, and 'signal' characterization and visualization. An important general conclusion of our study is the utility of a frequency-based representation, with short, moving analysis windows to account for nonstationarity in the data. Of particular note are (1) the development of a decomposition technique (space-frequency singular value decomposition) that is shown to be a useful means of characterizing the image data, and (2) the development of an algorithm, based on multitaper methods, for the removal of approximately periodic physiological artifacts arising from cardiac and respiratory sources.

In this chapter, we treat in succession dynamic brain imaging data gathered in MEG, optical, and fMRI experiments. Based on the analysis of actual data sets using the techniques introduced in Chapter 1, we discuss protocols for analysis to remove artifactual components and determine the structure of the signal.

2.2 Different brain imaging techniques

2.2.1 Imaging techniques and their spatiotemporal resolution

The three main techniques of interest here are optical imaging, fMRI, and MEG. Optical imaging falls into the further subcategories of intrinsic and extrinsic contrast. In extrinsic contrast optical imaging, an optical contrast agent sensitive to neuronal activity is added to the preparation. Examples of such contrast agents include voltage-sensitive dyes and Ca^{2+} concentration-sensitive contrast agents. Voltage-sensitive dye molecules inserted in the cell membrane and the small Stark shifts produced in the molecule by changes in the transmembrane voltage are sources of the contrast. The signal-to-noise ratio (SNR) in these experiments is typically poor, being of the order of unity. The spatial resolution is set by the optical resolution and scattering properties of the medium and can be of the order of microns. The temporal resolution is limited by the digitization rate of the recording apparatus (CCD camera or photodiode array) and can currently go up to ~ 1 kHz, which is the intrinsic timescale of neuronal activity. In calcium ion-sensitive imaging, the intrinsic timescales are slower, so the demands on the digitization rate are somewhat less. The SNR ratio is significantly better compared to currently available voltage-sensitive dyes. The spatial resolution is greatly enhanced in confocal (Pawley, 1995) and multiphoton scanning optical imaging (Denk et al., 1990). The imaging rates in multiphoton scanning optical imaging are currently significantly slower than the corresponding rates for CCD cameras.

Intrinsic optical imaging and fMRI rely on the same underlying mechanism, namely hemodynamic changes triggered by neuronal activity. Hemodynamic changes include changes in blood flow and blood oxygenation level. The intrinsic timescale for these changes is slow, ranging from hundreds of milliseconds to several seconds. The intrinsic spatial scale is also somewhat large, ranging from hundreds of microns to millimeters. These scales are well within the scope of optical techniques. In both the extrinsic and intrinsic cases discussed above, various noise sources, including physiological fluctuations, are important indirect determinants of the spatiotemporal resolution.

In fMRI, the instrumental limitations on the spatial and temporal resolutions are sig-

nificant, and for fixed SNR, a tradeoff exists between spatial and temporal resolution, as well as between temporal resolution and spatial coverage. The fMRI images are typically gathered in two-dimensional slices of finite thickness, and for a fixed SNR, the number of slices is roughly linear in time. For single slice experiments, the temporal resolution is ~ 100 ms, and the spatial resolution is ~ 1 mm. This spatial resolution can be improved for a single slice by sacrificing temporal resolution. For multiple slices covering the whole head, the temporal resolution is ~ 2 s. Note that these numbers may be expected to change somewhat based on future improvements in instrumentation. The principal advantage of fMRI is that it is noninvasive imaging, making it suitable for the study of the human brain. In addition, optical imaging is limited to the surface of the sample, whereas fMRI is a volumetric imaging technique.

In MEG, weak magnetic fields of the order of tens of fT/\sqrt{Hz} generated by electric currents in the brain are measured using superconducting quantum interferometric detector (SQUID) arrays positioned on the skull. Like fMRI, MEG is a noninvasive imaging technique and therefore applicable to the human brain. The temporal resolution (~ 1 ms) is much higher than in fMRI, although the spatial resolution is in general significantly poorer. The spatial resolution of MEG remains a debatable issue, because of the ill-posed nature of the inverse problem that must be solved in order to obtain an image from the MEG data. A promising direction for future research appears to be a combined use of fMRI and MEG performed separately on the same subject.

2.2.2 Sources of noise

As mentioned previously, the 'noise' present in the imaging data arise from two broad categories of sources, biological and nonbiological. Biological noise sources include cardiac and respiratory cycles, as well as motion of the experimental subject. In imaging studies involving hemodynamics such as fMRI and intrinsic optical imaging, an additional physiological source of noise is slow 'vasomotor' oscillations (Mitra et al., 1997; Mayhew et al., 1996). In addition in all studies of evoked activity, ongoing brain activity not locked to or triggered by the stimulus appears as noise. Nonbiological noise sources include photon counting noise in

optical imaging experiments, noise in the electronic instrumentation, 60 Hz noise, building vibrations, and the like.

We first consider optical imaging using voltage-sensitive dyes in animal preparations. The sensitivity of these experiments is currently limited by photon counting noise. The Stark shifts associated with the available dyes lead to changes in the optical fluorescence signal on the order of $\Delta F/F \sim 10^{-3}/mV$. The typical SNR is therefore of order unity or less. In addition, absorption changes arising from hemodynamic sources, whether related to the stimulus or not, corrupt the voltage-sensitive dye images. Perfusing the brain with an artificial oxygen-supplying fluid can eliminate these artifacts (Prechtl et al., 1997). Motion artifacts and electronic noise may also be significant. In contrast, calcium-sensitive dye images have comparatively large signal changes for spike mediated Ca^{2+} fluxes, and are less severely affected by the photon counting noise. However, motion artifacts can still be severe, particularly at higher spatial resolutions.

In fMRI experiments in humans, the instrumental noise is small, typically a fraction of a percent. The dominant noise sources are of physiological origin, mainly cardiac, respiratory and vasomotor sources (Mittra et al., 1997). Depending on the time between successive images, these oscillations may be well resolved in the frequency spectrum, or aliased and smeared over the sampling frequency interval. Subject motion, resulting from respiration or other causes, is a major confounding factor in these experiments.

In MEG experiments the SNR is usually very low. Hundreds of trial averages are sometimes needed to extract evoked responses. Magnetic fields from currents associated with the cardiac cycle are a strong noise source, as are 60 Hz electrical sources and other nonbiological current sources. In studying the evoked response, the spontaneous activity not related to the stimulus is a dominant source of undesirable fluctuations.

2.3 Description of data sets

In this section we describe the data sets used to illustrate the techniques developed in the paper. We have used data from three different imaging techniques, corresponding to

multichannel MEG recordings, optical image time series, and MRI time series. The data are grouped into four sets, referred to below using the letters \mathcal{A} through \mathcal{D} . Brief discussions of the data collection are provided below. The data have either been previously reported or are gathered using the same techniques as in other reports. In each case, a more detailed description may be found in the accompanying reference.

Data set \mathcal{A} consists of multichannel MEG recordings. Descriptions of the apparatus and experimental methods for these data can be found in (Joliot et al., 1994). For our purposes it is sufficient to note that the data were gathered simultaneously from 74 channels using a digitization rate of 2.083 kHz for a total duration of 5 min, and correspond to magnetic fields resulting from spontaneous brain activity recorded from an awake human subject in a resting state with eyes closed.

Data set \mathcal{B} consists of dynamic optical images of the procerebral lobe of the terrestrial mollusc *Limax* (Kleinfeld et al., 1994), gathered after staining the lobe with a voltage-sensitive dye. The digitization rate is 75 Hz and the total duration of the recording is 23 s. The images are 105×34 pixels in extent and cover an area of $\sim 600\mu m \times 200\mu m$.

Data sets \mathcal{C} and \mathcal{D} contain fMRI data and consist of time series of magnetic resonance images of the human brain showing a coronal slice toward the occipital pole (Mitra et al., 1997; Le and Hu, 1996). The data were gathered in the presence (\mathcal{C}) or absence (\mathcal{D}) of a visual stimulus. For data set \mathcal{C} , binocular visual stimulus was provided by a pair of flickering red LED patterns (8 Hz), presented for 30 s starting 40 s after the beginning of image acquisition. The digitization rate for the images was 5 Hz and the total duration 110 s. The images are 64×64 pixels and cover a field of view of $20cm \times 20cm$.

2.4 Analysis techniques for different modalities

We now deal with specific examples of data from different brain imaging modalities. We concentrate on general strategies for dealing with such data, and most of our techniques apply with minor changes across the modes. However, it is easier to treat the various cases separately, and we present somewhat parallel developments in the following sections. The

data sets \mathcal{A} through \mathcal{D} have been briefly discussed before.

2.4.1 Magnetoencephalography

In this section we consider data set \mathcal{A} consisting of multichannel MEG recordings. The number of channels is 74 (37 on each hemisphere), the digitization rate 2.083 kHz, and the duration of the recording is 5 min.

2.4.1.1 Preprocessing

An occasional artifact in MEG recordings consists of regularly spaced spikes in the recordings due to cardiac activity. This is caused when the comparatively strong magnetic field due to currents in the heart is not well canceled. To suppress this artifact, we proceed as follows: first, a space-time SVD is performed on the multichannel data. The cardiac artifacts are usually quite coherent in space, and show up in a few principal component time series. In those time series, the spikes are segmented out by determining a threshold by eye and segmenting out 62.5 ms before and after the threshold crossing. The segmented heartbeat events are then averaged to determine a mean waveform. Since the heartbeat amplitude is not constant across events, the heartbeat spikes are removed individually by fitting a scaling amplitude to the mean waveform using a least-squares technique. Each spike modeled by the mean waveform multiplied by a scaling amplitude is then subtracted from the time series. Figure 2.1 illustrates the results of this procedure.

A fairly common problem in electrical recordings is the presence of 60 Hz artifacts and on occasion sinusoidal artifacts at other frequencies. Such sinusoidal artifacts, if they lie in the relevant data range, are usually dealt with using notch filters. This is, however, unnecessarily severe since the notch filters may remove too large a band of frequencies, in particular in the case of MEG data where frequencies close to 60 Hz are of interest. We find that the 60 Hz line and other fixed sinusoidal artifacts can be efficiently estimated and removed using the methods for sinusoidal estimation described in Section 1.1.5. Specifically, in this case the frequency of the line is accurately known (this requires a precise knowledge of the digitization rate) and one has only to estimate its amplitude and phase. This is done

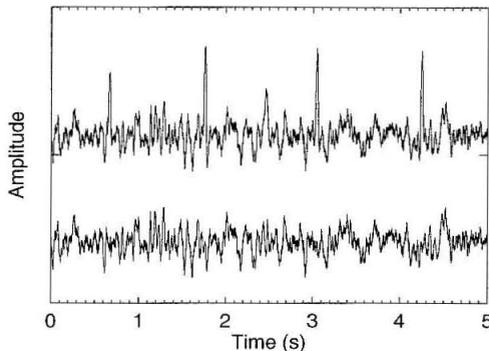


Figure 2.1: PC time series of MEG data from data set \mathcal{A} containing cardiac artifacts. The upper plot shows the time course before suppression and the spikes due to the heart beat are clearly visible. The lower plot shows the time course after suppression of the heart beat using the technique described in the text.

for a small time window using Eq.1.12. By sliding this time window along, one obtains a slowly varying estimate of the amplitude and phase, and is therefore able to reconstruct and subtract the sinusoidal artifact. The results of such a procedure are illustrated in Figure 2.2a, where a time averaged spectrum is shown for a single channel before and after subtraction of the line frequency component. In Figure 2.2b, the amplitude of the estimated 60 Hz component is shown with its slow modulation over time.

2.4.1.2 Time-frequency analysis

The spectral analysis of multichannel data presents the fundamental problem of how to simultaneously visualize or otherwise examine the time-frequency content of many channels. One way to reduce the dimensionality of the problem is to work with PC time series. A space-time SVD of MEG data gives a rapidly decaying singular value spectrum. This indicates that one can consider only the first few temporal components to understand the spectral content of the data. For purposes of illustration of our techniques we consider the first three PC time series in a 5 min segment recording of spontaneous awake activity.

A useful preliminary step in spectral analysis is prewhitening using an appropriate autoregressive model. This is necessary because the spectrum has a large dynamic range.

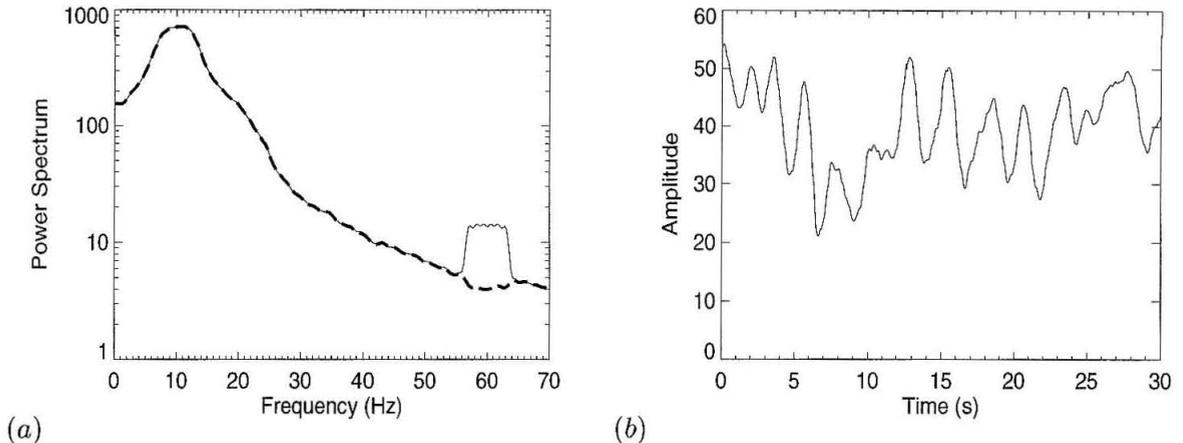


Figure 2.2: (a) Time averaged spectrum before (*solid line*) and after (*dotted line*) subtraction of 60 Hz artifact (single channel time series) from MEG data in data set \mathcal{A} . (b) Estimated time varying amplitude of the 60 Hz line frequency.

Prewhitening leads to equalization of power across frequencies, which allows better visualization of time-frequency spectra. In addition, a space-time SVD is better performed on prewhitened data, since otherwise the large amplitude, low frequency oscillations completely dominate the principal components. The qualitative character of the average spectrum is a slope on a semilogarithmic scale (Fig. 2.2a). The goal is to prewhiten with a low order AR model so that the peaks in the spectrum are left in place, but the overall slope is removed. Considering the derivative of the spectrum rather than the spectrum itself achieves a similar result.

The procedure is to first calculate a moving estimate of the spectrum using a short time window ($T = 0.35s$ in the present case) and a direct multitaper spectral estimator ($W = 4/T; K = 6$). These estimates are then averaged over time to obtain a smooth overall spectrum. Next, a low order autoregressive model (order = 10 in the present case) is fit to the spectral estimate. We use the Levinson Durbin recursion to fit the AR model. Results of such a fit are shown in Figure 2.3. The coefficients of the autoregressive process are then used to filter the data, and the residuals are subjected to further analysis. Thus, if the coefficients obtained are a_k and the original time series is x_t , then the residuals are

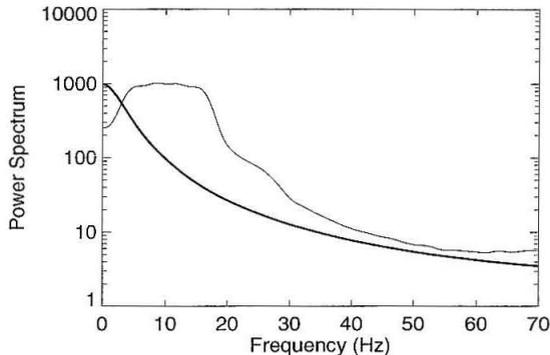


Figure 2.3: AR fit to averaged spectrum of MEG data from data set \mathcal{A} averaged over non-overlapping windows in time for purposes of prewhitening. The thick line shows the autoregressive fit to the average spectrum. The thin line shows the average spectrum. A low-order AR spectral estimate is used to reduce the dynamic range in the spectrum without fitting specific structural features of the data.

$\delta x_t = x_t - \sum_1^n x_{t-k} a_k$, which are subjected to a time-frequency analysis.

Typical time-frequency spectra of prewhitened PC time series are shown in Figure 2.4. The spectra were obtained using a direct multitaper estimate for 1s long time windows and for $W = 4\text{Hz}$, $K = 6$.

To assess the quality of the spectral characterization it is important to quantify the presence (or absence) of correlations between fluctuations at different frequencies. One measure of the correlations between frequencies for a given time series is given by the following quantity analogous to the coherence between different channels:

$$\rho_{xx}(f, f') = \frac{\sum_{k=1}^K \tilde{x}_k(f) \tilde{x}_k^*(f')}{\sqrt{\sum_{k=1}^K |\tilde{x}_k(f)|^2 \sum_{k=1}^K |\tilde{x}_k(f')|^2}} \quad (2.1)$$

The estimate can be further averaged across time windows to increase the number of degrees of freedom. This quantity, computed for the leading PC time series obtained earlier, is displayed in Figure 2.5a. In this figure, the magnitude of the estimate $\rho(f, f')$ for the leading PC time series is displayed as a function of the two arguments, f and f' . For comparison, in Figure 2.5b, results of the same procedure obtained after initially

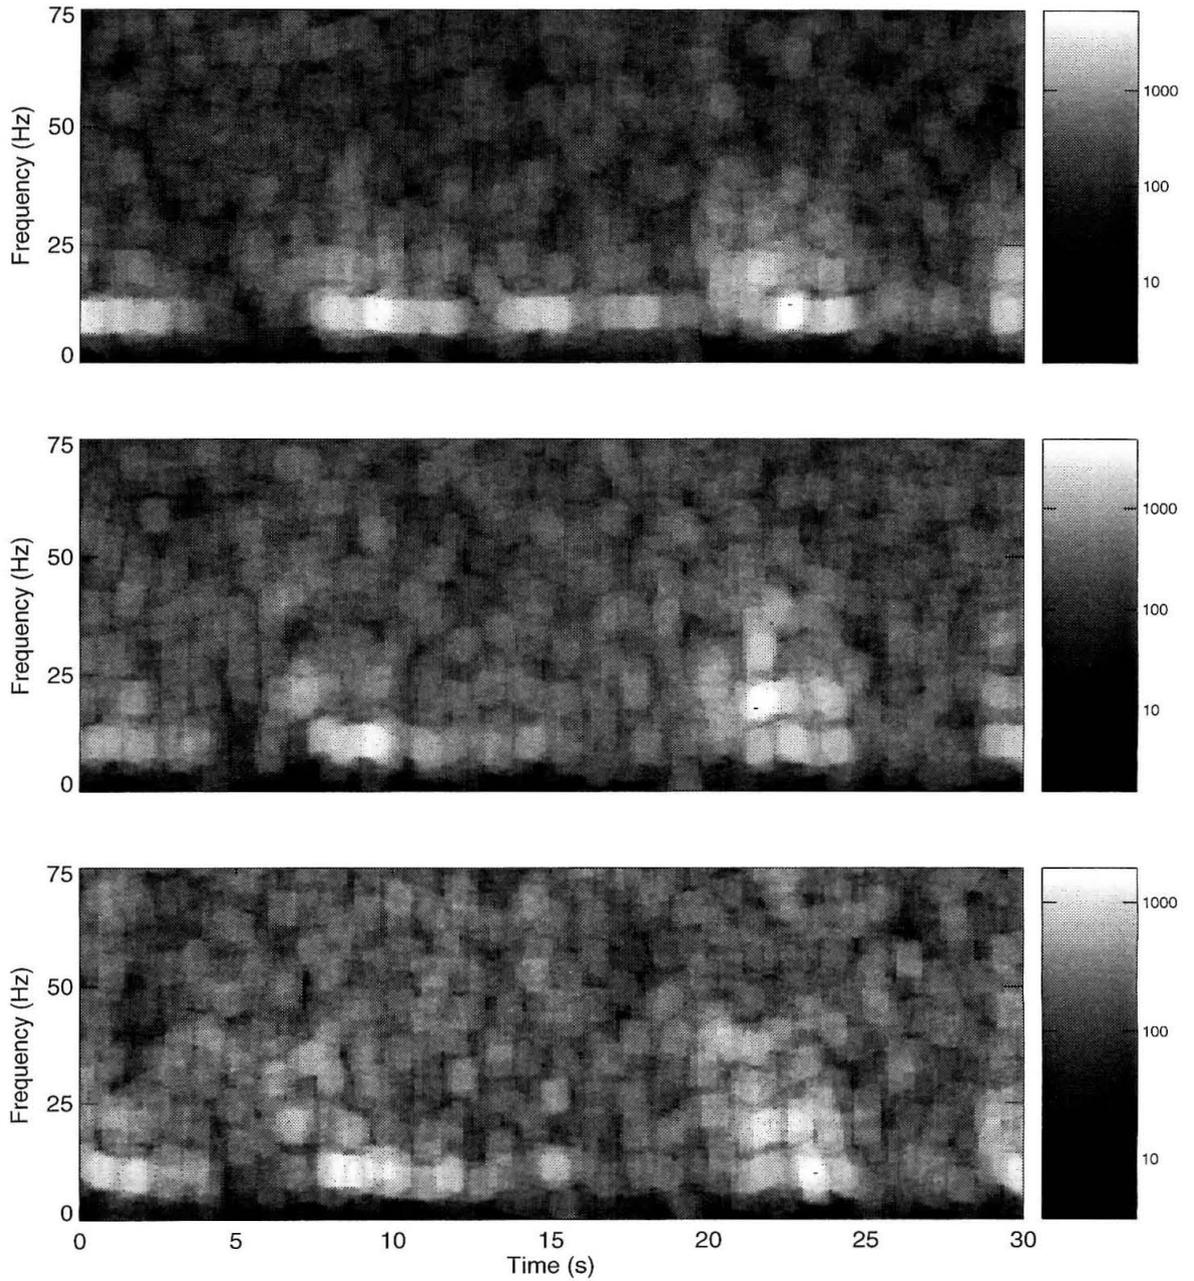


Figure 2.4: Time frequency spectrum of first three PCs obtained in a space-time SVD of MEG data from data set \mathcal{A} pre-whitened by the filter shown in Figure 2.3.

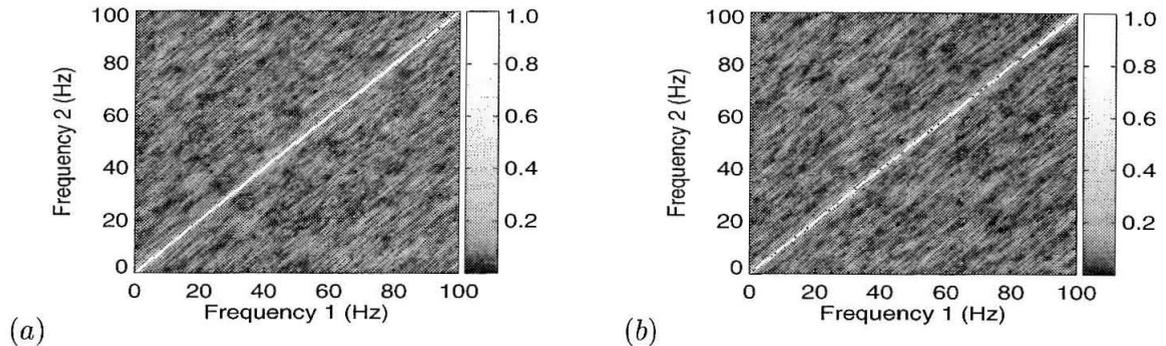


Figure 2.5: (a) Magnitude of $\rho(f, f')$ for the leading PC time series of MEG data from data set \mathcal{A} . (b) Magnitude of $\rho(f, f')$ for the leading PC time series after initially scrambling the time series.

scrambling the time series are also displayed. A visual comparison shows the lack of evidence for correlations across frequencies in an average sense. This rules out, for example, the scenario in which transient peaks in the spectrum at high frequencies are predominantly harmonics of corresponding transient peaks at low frequencies. Although in the present case we obtain a null result, it can be expected that this measure will provide useful information when correlations between different frequencies are actually present.

2.4.1.3 Multichannel spectral analysis

The time frequency spectra of leading PCs shown in the earlier section capture the spectral content of the MEG signal that is coherent in space. Alternatively, one can perform a space-frequency SVD with a moving time window on the data. It is also desirable to obtain time-averaged characterizations of the coherence across channels. This can be done by considering the coherence functions between channels $\rho_{ij}(f)$, which can be estimated as in Eq.1.26 using multitaper methods. The estimate, when calculated for a moving time window, can be further averaged across time windows. Displaying the matrix ρ_{ij} poses a visualization problem, since the indices i, j themselves correspond to locations on a two-

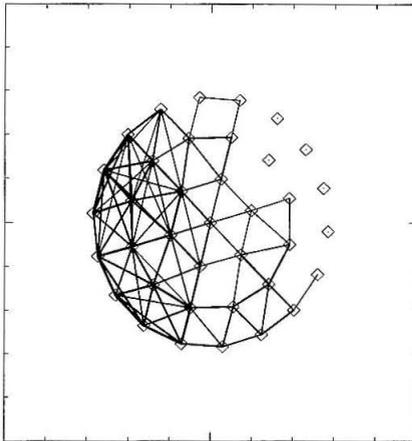


Figure 2.6: Coherence of MEG activity from data set \mathcal{A} between points in space displayed by lines of proportionate thickness. Coherence was computed for a center frequency of 20 Hz and half bandwidth 5 Hz.

dimensional grid. Thus, an image displaying the matrix ρ_{ij} does not preserve the spatial relationships between channels. One solution to this visualization problem is presented in Figure 2.6, by representing the strength of the coherence between two space points by the thickness of a bond connecting the two points. This figure shows the coherence $\rho_{ij}(f)$ computed with a center frequency 20 Hz and half bandwidth 5 Hz. The bond strengths have been thresholded to facilitate the display. This visualization, although not quantitative, allows for an assessment of the organization of the coherences in space.

An alternative way of performing principal component analysis on space-time data while localizing information in the frequency domain is clearly to apply a space-time SVD to data which has first been frequency filtered into the desired band. To obtain frequency filters with optimal band-limiting properties, projection filters based on DPSS are used. For illustration, we perform this analysis on the data under discussion. The individual channels were first filtered into the frequency band 35–45 Hz. This gives a complex time series at each spatial location. The first three dominant spatial eigenmodes in a space-frequency SVD are displayed in Figure 2.7, with singular values decreasing from the top to the bottom of the

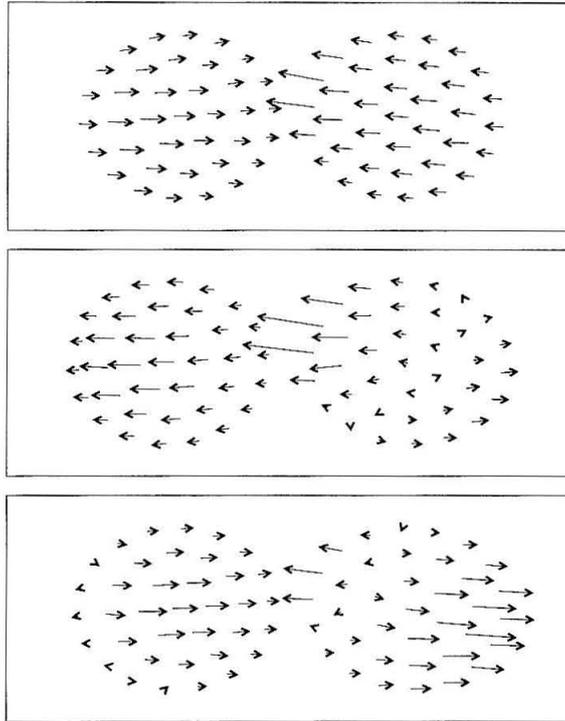


Figure 2.7: Dominant spatial eigenmodes of a space-time SVD of band pass filtered MEG data from data set \mathcal{A} for the frequency band 35–45 Hz. The two hemispheres of the brain are projected onto a plane, each being sampled by 37 sensors. There is clearly a high degree of spatial coherence at this frequency band.

figure. Since the spatial eigenmodes are complex, their values are represented by arrows, whose lengths correspond to magnitudes and whose directions correspond to phases. It is quite clear from the figure that the data show a high degree of spatial coherence on an average.

2.4.2 Optical imaging

In this subsection we consider optical imaging data. We consider the general case of imaging data gathered either using intrinsic or extrinsic contrast. The case of intrinsic contrast is closely related to fMRI, and the analysis parallels that of the fMRI data sets \mathcal{C} and \mathcal{D} . To illustrate some effects important in the case of extrinsic contrast, we consider data set \mathcal{B} . Data set \mathcal{B} was gathered in presence of a voltage sensitive dye, and consists of images of an isolated procerebral lobe of *Limax*, with a digitization rate 75 Hz and duration 23 s.

2.4.2.1 Spectral analysis of PC time series

The general procedure outlined above for fMRI data consisting of a space-time SVD followed by a spectral analysis of the PC time series is useful to obtain a preliminary characterization of the data. In particular, in the case of optical imaging data using intrinsic or extrinsic contrast in the presence of respiratory and cardiac artifacts, this procedure helps the assessment of the artifactual content of the data. However, as discussed earlier, the space-time SVD mixes up distinct dynamic components of the image data, and is therefore of limited utility for a full characterization of the data.

2.4.2.2 Removal of physiological artifacts

We have developed a method for efficient suppression of respiratory and cardiac artifacts from brain imaging data (including optical images and MRI images for high digitization rates) by modeling these processes by slow amplitude- and frequency-modulated sinusoids. Other approaches to suppression of these artifacts in the literature include “gating,” frequency filtering, modeling of the oscillations by a periodic function (Le and Hu, 1996),

and removal of selected components in a space-time SVD (Orbach et al., 1995). These approaches have varying efficacies. For example, gating aliases the relevant oscillations down to zero frequency, so that any variation in these oscillations cause slow fluctuations in the data, which is in general undesirable. Frequency filtering removes more spectral energy than is strictly necessary from the signal. Modeling of the oscillations by a periodic function is imperfect because the oscillations themselves may vary in time. This can be rectified by allowing the parameters of the oscillations to slowly change in time. One must be able to fit a sinusoidal model robustly to short time series segments to do this properly. Finally, removing selected components in a space-time SVD is not a safe procedure because, as discussed before, the space-time SVD does not necessarily separate the different components of the image data.

Our method for suppressing the above mentioned oscillatory components is based on multitaper methods for estimating sinusoids in a colored background described earlier. The method is based on modeling the oscillations by a sum of sinusoids whose amplitude and frequency are allowed to vary slowly. The modeled oscillations are removed from the time series in the data to obtain the desired residuals. Although a space-time SVD is not sufficient by itself, applying such a removal technique to the leading temporal PCs, and reconstituting the residual time series appears to give good results. This procedure was found to be effective for a wide variety of data, including optical imaging using both intrinsic and extrinsic contrast in rat brain, and in fMRI data. The details of the technique are presented in the section on fMRI data.

2.4.2.3 Space-frequency SVD

In this section we present the results of a space-frequency SVD applied to data set \mathcal{B} . This technique has been described above (Figs. 1.9, 1.10) for data set \mathcal{C} . In Figure 2.8 the coherence spectrum is shown for data set \mathcal{B} . The coherence was computed on a coarse grid, since there is not much finer structure than that displayed in the spectrum. In this case, 13 DPSS were used, corresponding to a full bandwidth of 0.3 Hz. The preparation (procerebral lobe of Limax) is known to show oscillations, which are organized in space as a

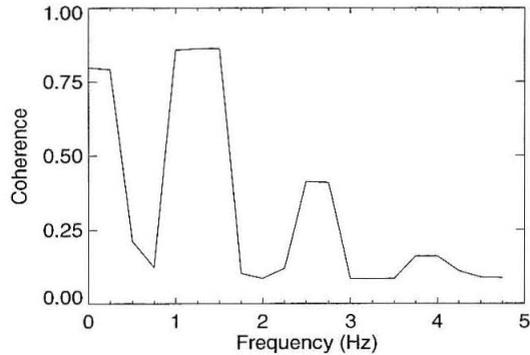


Figure 2.8: Coherence spectrum for a space-frequency SVD of optical imaging data in data set \mathcal{B} from the procerebral lobe of *Limax*.

traveling wave. Traveling waves in the image, gathered in the presence of a voltage sensitive dye, reflect traveling waves in the electrical activity. The coherence spectrum displayed in Figure 2.8 shows a fundamental frequency of about 1.25 Hz and the corresponding first two harmonics.

The amplitudes of the leading spatial eigenmode as a function of center frequency are shown in Figure 2.9. Note that the spatial distribution of coherence is more localized to the center of the image at the higher harmonics. This reflects the change in shape of the waveform of oscillation that is known to occur in the preparation as a function of spatial position. This phenomenon has been interpreted as a result of differing spatial concentration profiles of two different cell types in this system (Kleinfeld et al., 1994). Based on this past interpretation, the leading modes at the fundamental and harmonic frequencies directly reflect the spatial distribution of the different cell types.

The spatial eigenmodes are complex and possess a phase in addition to an amplitude. This allows for the investigation of traveling waves in the data. Let the leading spatial eigenmode, as a function of frequency, be expressed as

$$\tilde{I}_1(x; f) = A(x; f) \exp(i\theta(x; f)) \quad (2.2)$$

Given the convention we are following for the Fourier transform (Eq.1.1), one may define

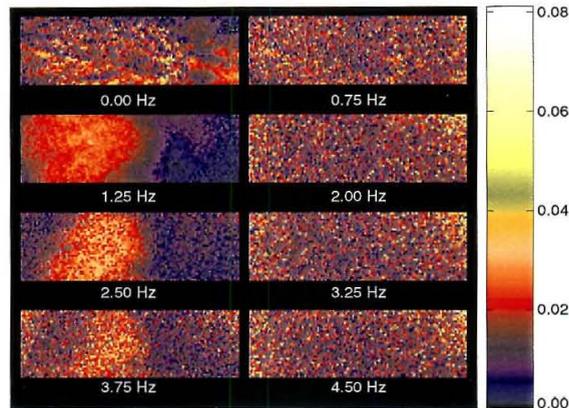


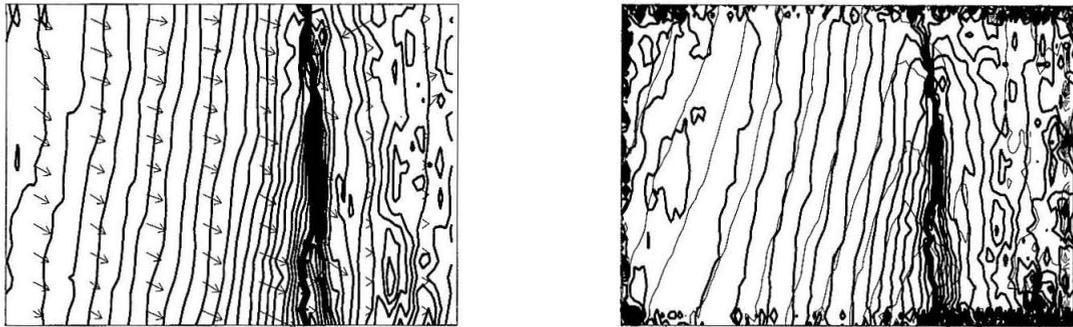
Figure 2.9: Leading spatial eigenmodes of the space-frequency SVD corresponding to Figure 2.8. The amplitudes of the leading modes are shown as a function of center frequency.

the following local wave vector:

$$\mathbf{k}(x; f) = -\nabla\theta(x; f) \quad (2.3)$$

If the coherent fluctuations at a given frequency correspond to traveling plane waves, then $\mathbf{k}(x; f)$ corresponds to the usual definition of the wave vector. More generally, this quantity allows the systematic examination of phase gradients in the system, which corresponds to traveling excitations.

The local wave vector map for data set \mathcal{B} at a center frequency of 2.5 Hz (the first harmonic) is shown in Figure 2.10a, superposed on top of contours of constant phase. In Figure 1.10b, the constant phase contours for center frequencies 1.25 Hz, 2.5 Hz (fundamental and first harmonic) are superposed. On superposing the contours for the fundamental and the first harmonic, we discover an effect that was not evident in the earlier analysis of the data (Kleinfeld et al., 1994), namely that the phase gradients at 1.25 and 2.5 Hz are slightly tilted with respect to each other. This can be interpreted as two different waves



(a)

(b)

Figure 2.10: Leading spatial eigenmodes of the space-frequency SVD corresponding to Figure 2.9. (a) Gradients of the phase of the leading SVD mode at a center frequency of 2.5 Hz, corresponding to local wave-vectors for the wave motion, are shown as arrows. These are superposed on constant phase contours for the mode. (b) Constant phase contours for spatial eigenmodes of the space-frequency SVD for center frequencies 1.25 Hz and 2.5 Hz. 2.5 Hz is shown by the bold contour.

simultaneously present in the system, but running in slightly different directions. Coexisting waves present at different temporal frequencies with different directions of propagation have also been revealed by space-frequency SVD analysis of voltage-sensitive dye images of turtle visual cortex (Prechtl et al., 1997).

2.4.3 Magnetic resonance imaging

The data sets \mathcal{C} and \mathcal{D} comprising functional MRI data have been used in earlier sections to illustrate the techniques presented in the paper (Figs. 1.7–1.10). In this section we continue to illustrate analytical techniques on this data set. The data sets were gathered with a digitization rate of 5 Hz and a total duration of 110 s. Data set \mathcal{C}/\mathcal{D} was gathered in the presence/absence of a flashing LED checkerboard pattern serving as visual stimulus. An extra problem in the analysis of MRI data is the presence of motion related artifacts, which have to be suppressed (Mittra et al., 1997).

2.4.3.1 Removal of physiological artifacts

Here a detailed description is provided of the method for removal of physiological oscillations discussed in the section on optical imaging data. A space-time SVD of the data is first computed, followed by sinusoidal modelling of the leading principal component time series. This is necessary for two reasons: (1) The images in question typically have many pixels, and it is impractical to perform the analysis separately on all pixels. (2) the leading SVD modes capture a large degree of global coherence in the oscillations.

Consider a single principal component time series, $a(t)$. We assume that the time series is a sum of two components. The first component consists of a sum of amplitude- and frequency-modulated sinusoids representing respiratory and cardiac oscillations. The second component $\delta a(t)$ contains the desired signal.

$$a(t) = \sum_n A_n(t) \cos[f_n(t)t + \phi_n(t)] + \delta a(t) \quad (2.4)$$

The goal is to estimate the smooth functions $A_n(t)$, $f_n(t)$, and $\phi_n(t)$, which give the component to be subtracted from the original time series.

It is necessary to choose an optimally sized analysis window. This window must be sufficiently small to capture the variations in the amplitude, frequency, and phase, but must be long enough to have the frequency resolution to separate the relevant peaks in the spectrum, both artifactual and originating in the desired signal. The choice of window size depends to some extent on the nature of the data, and cannot be easily automated. However, in similar experiments it is safe to use the same parameters. Ideally, one would choose the window size in some adaptive manner, but we find it adequate for our present purposes to work with a fixed window size.

The frequencies $f_n(t)$ have to satisfy several criteria. Usually one is removing the respiratory and cardiac components. The corresponding spectra contain small integer multiples of two fundamental frequencies, with the possible presence of sidebands due to nonlinear interactions between the oscillations. The frequency F-test described in Multitaper spectral analysis is used to determine the fundamental frequency tracks $f_n(t)$ in Equation 2.4. The

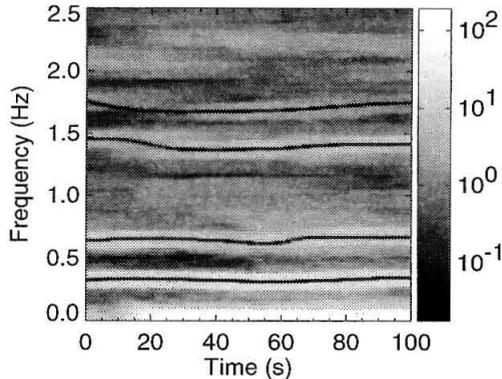


Figure 2.11: Time frequency plot of frequency tracks for a moving sinusoidal model of the cardiac and respiratory artifacts in fMRI data from data set \mathcal{C} . The analysis is performed on a PC time series obtained by performing a space-time SVD.

time series used for this purpose may either be a PC time series or an independently monitored physiological time series. Note that the assumption here is that within an analysis bandwidth of the relevant peak in the cardiac or respiratory cycle, the data can be modeled as a sine wave in a locally white background. The fundamental frequency tracks are used to construct the tracks for the harmonics and the sidebands.

In the example, we show the results of the analysis on one principal component time series from data set \mathcal{C} . The fundamental frequency tracks were determined by using the F-test on a moving analysis window on the PC time series.

If the F-test does not provide a frequency estimate for short segments of the data, estimates may be interpolated using a spline, for example. After the frequency tracks are determined, the amplitude and phase of the sinusoids are calculated using Equation 1.12 for each analysis window location. Note that the shift in time between two successive analysis windows can be as small as the digitization rate of the data, but is limited in practice by the available computational resources. The estimated sinusoids are reconstructed for each

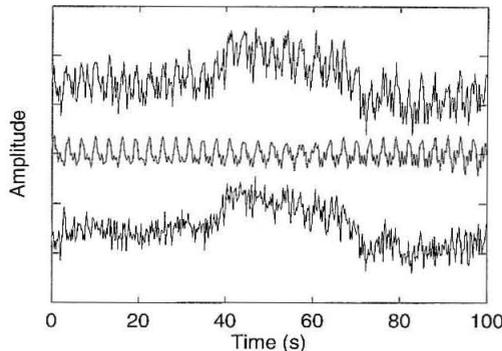


Figure 2.12: *Top:* Part of PC time series forming the basis of Fig.23. *Middle:* Estimated cardiac and respiratory components corresponding to top plot. *Bottom:* Time series with cardiac and respiratory components suppressed.

analysis window, and the successive estimates are overlap-added to provide the final model waveform for the artifacts.

The frequency tracks are shown in Figure 2.11 superposed on a time-frequency spectral estimate of the principal component time series. In Figure 2.12, results of the procedure described above are shown in the time domain for the chosen PC. Notice the change in the frequency corresponding to the cardiac cycle (~ 1.3 Hz) in the initial part of the time period. This would prevent the adequate estimation of this component if a model with fixed periodicity were used. In contrast the method used here allows for slow variations in the amplitudes of the oscillations in addition to variations in frequency. A strong stimulus response is noticeable in this PC (data set \mathcal{C} was gathered in the presence of visual stimulus).

2.4.3.2 Space-frequency SVD

In this section, the results of a space-frequency SVD are shown for data sets \mathcal{C} and \mathcal{D} . Recall that data sets \mathcal{C} and \mathcal{D} were collected with identical protocols, except that in data set \mathcal{C} a controlled visual stimulus is applied. Figure 2.13 shows the coherence spectra resulting from the space-frequency SVD. In this calculation, the DPSS used corresponded to a full bandwidth of 0.1 Hz. The coherence spectra for the two data sets are more or less

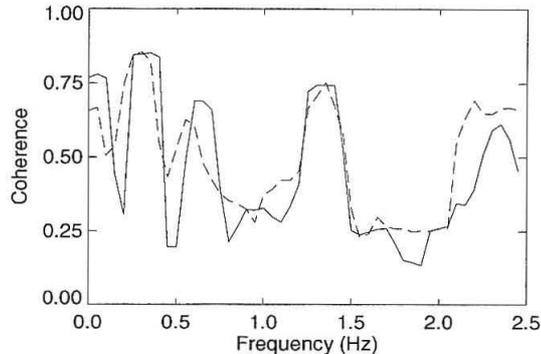


Figure 2.13: Coherence spectrum for a space-frequency SVD of fMRI data from data sets \mathcal{C} and \mathcal{D} in presence of visual stimulus (*solid curve*) compared with absence of visual stimulus (*dotted curve*).

the same. The coherence near zero frequency is higher for data set \mathcal{C} which contains the visual stimulus. The stimulus response can be seen clearly in the amplitude of the leading spatial eigenmode of the space-frequency SVD for data set \mathcal{C} (Fig. 2.14) at close to zero frequency. At higher frequencies, coherence arising from artifactual (respiratory) sources causes a different pattern of spatial amplitudes. As opposed to the space-time SVD, this procedure segregates the stimulus response from the oscillatory artifacts.

2.5 Discussion

We have tried to outline analysis protocols for data from the different modalities of brain imaging. It is useful to recapitulate the essential features of these protocols in a unified manner, and to indicate the domains of validity of the different techniques proposed.

Visualization of raw data

It is usually necessary to directly visualize the raw data, both as a crude check on the quality of the experiment and to direct further analysis. In this stage, one may look at individual time series from the images or look at the data displayed dynamically as a movie. The relevant images are often noisy, so a noise reduction step is first necessary even before the preliminary visualization. In cases where the visualization is limited by large shot noise,

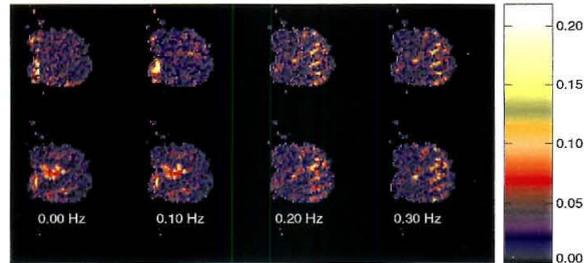


Figure 2.14: Amplitudes of leading spatial eigenmodes corresponding to the space-frequency SVD in Figure 2.13. Center frequencies from 0 Hz to 0.3 Hz in 0.1 Hz increments. *Top:* Without stimulus. *Bottom:* With stimulus.

truncation of a space-time SVD with possibly some additional smoothing provides a simple noise reduction step for the visualization.

Preliminary characterization

In the next stage, it is useful to obtain quantities that help parse out the content of the data, in particular to identify the various artifacts. Despite its limitations, a space-time SVD is useful at this stage to reduce the data to a few time series and corresponding eigenimages. Examination of the aggregate spectra of the PC time series, for example, reveals the extent of cardiac/respiratory content of fMRI/optical imaging data. In case of MEG, direct examination of the PC time series reveals the degree of cardiac contamination. Examination of the corresponding spatial images reveals the spatial locations of the artifacts. In case of fMRI data, where the digitization rate may not be very high, studying the spectra can reveal whether cardiac/respiratory artifacts still lead to possibly aliased frequency peaks in the power spectra.

A further, more powerful characterization is obtained by the space-frequency SVD. For optical data and for rapidly sampled fMRI data, there is sufficient frequency resolution that at this stage the oscillatory artifacts segregate well. Studying the overall coherence

spectrum reveals the degree to which the images are dominated by the respective artifacts at the artifact frequencies, while the corresponding leading eigenimages show the spatial distribution of these artifacts more cleanly compared to the space-time SVD. Moreover, provided the stimulus response does not completely overlap the artifact frequencies, a characterization is also obtained of the spatiotemporal distribution of the stimulus response. In case of fMRI, if the digitization rates are too slow (< 0.3 Hz), there may not be any segregation in the frequency domain of the various components of the image; this can be established at this stage by examining the eigenimages of the space-frequency SVD. In this case, the techniques described in this paper would be of limited use.

Artifact removal

Based on the preliminary inspection stage, one can proceed to remove the various artifacts to the extent possible. The techniques described in this paper are most relevant to artifacts that are sufficiently periodic, such as cardiac/respiratory artifacts in optical/fMRI data, 60 Hz noise in optical data/MEG, and other frequency-localized noise such as building/fan vibrations (optical imaging data). There are two basic ways of using the frequency segregation of the artifacts to remove them. One method is to directly model the waveforms of the oscillations using the frequency- and amplitude-modulated sinusoidal fit described in the sections on optical/fMRI data.

For fMRI data, if the digitization rate is too low, then the techniques described here are not useful. However, for fMRI data, even with digitization rates of 1–2 s, it appears possible to use the frequency segregation of the physiological artifacts, using one of two methods. If auxiliary time series are available for cardiac/respiratory oscillations, one may construct the transfer function from these time series to the data using the multitaper technique described above, perform a statistical test of significance such as the F-test, and remove the significantly fitted components. Alternatively, if no auxiliary data are available, the space-frequency SVD may be examined for presence of these artifacts, and if the artifacts can be identified with some frequency band then filtering techniques may be used.

In each of the cases described above, the fundamental operation is performed on an individual time series. This may be performed pixel by pixel in the image, or to reduce

computational time, the steps may alternatively be performed on the leading PC time series and the artifacts thus reconstructed may be then subtracted from the raw data.

Stimulus response characterization

This may be the most delicate step, since the goal of the experiment is usually to find the stimulus response which is not known *a priori*. If the stimulus is presented periodically and repeatedly, the transfer function may be computed in the frequency domain using the techniques described before. Often, some strong assumption is made about the stimulus response (*e.g.*, the image intensity will rise during the stimulus), and methods of signal detection theory or statistical hypothesis testing are applied to extract the response based on the assumed signal model.

One general assumption about the stimulus response to a single trial might be that it lives in a particular region of frequency space. In the case of fMRI data, where the response is often a prolonged increase in signal intensity, this would correspond to the signal having relatively low frequencies. In this case, the space-frequency SVD described in the paper is of utility in describing the stimulus response, as illustrated in the section on fMRI data. Similarly, in optical imaging, if the stimulus was modulated with a particular temporal frequency, the same idea would be applicable.

It is to be emphasized that it is not possible to develop "black box"-like techniques that are a panacea to all problems of data analysis. It is neither possible nor desirable to entirely eliminate the human component in the process. However, relevant computational and analytical tools can be a powerful aid to making sense of the data, and are used most effectively in a closed loop system where the results of analysis influence experimentation.

Chapter 3 Temporal structure during working memory in Macaque parietal cortex

3.1 Introduction

Working memory is a brain system requiring the temporary storage and manipulation of information necessary for cognition (Baddeley, 1992). The neurophysiological basis of working memory is studied in non-human primates by recording activity during delayed response tasks (Fuster, 1995). Cue-selective elevations in single unit mean firing rates have been recorded during the delay period in many brain areas during different versions of these tasks (Fuster and Jervey, 1982; Bruce and Goldberg, 1985; Gnadt and Andersen, 1988; Miyashita and Chang, 1988; Funahashi et al., 1989; Koch and Fuster, 1989; Miller et al., 1996; Zhou and Fuster, 1996). Parietal cortex is important to spatial cognition (Andersen, 1995) and spatially-tuned increases in firing rate during working memory were first reported in parietal cortex on the lateral bank of the intraparietal sulcus (area LIP) (Gnadt and Andersen, 1988). In analogy to receptive fields in the sensory system, such activity can be said to form memory fields (Funahashi et al., 1989) and is thought to reflect the plan to make a movement (Mazzonni et al., 1996; Snyder et al., 1997; Batista and Andersen, 2001). Converging evidence indicates memory fields may be important to understanding the neural basis of working memory (Goldman-Rakic, 1995).

Temporally correlated neuronal activity in the form of reverberations has long been thought to be important to short-term memory (Lorente de No, 1938; Hebb, 1949; Amit, 1995) but measures of mean activity, such as the firing rate, do not capture correlations and measures of the variance, such as the spectrum, are required. If reverberant activity is present in neuronal activity, temporal structure localized in frequency may also predict planned saccades would be evidence for memory fields of temporal structure, or *dynamic*

memory fields. Dynamic memory fields are defined by cue-selective changes in the temporal structure of neuronal activity. The temporal structure of neuronal activity has been the subject of great interest (Singer and Gray, 1995; Roskies, 1999) and work has related it to perception (Eckhorn et al., 1988; Gray and Singer, 1989; Engel et al., 1990; Kreiter and Singer, 1996; Cardoso de Oliveira et al., 1997; Friedman-Hill et al., 2000), attention (Fries et al., 2001) and action (Rouguel et al., 1979; Sanes and Donoghue, 1993; Bressler et al., 1993; Murthy and Fetz, 1996a,b; Roelfsema et al., 1997; Donoghue et al., 1998; Lebedev and Wise, 2000). Previous studies of neural activity during object working memory have reported temporal structure in spike activity in monkeys (Nakamura et al., 1992) and in EEG activity in humans (Tallon-Baudry et al., 1999). The presence of temporal structure specific to a remembered movement plan would lend experimental support to the idea that working memory, and potentially other cognitive processes, involves reverberating neuronal activity.

We recorded spiking and local field potential (LFP) activity from two macaques during a memory-saccade task using a single tetrode located in area LIP. Spike activity has been previously examined in area LIP during this task (Gnadt and Andersen, 1988; Chafee and Goldman-Rakic, 1998). Previous analysis of this data has investigated it for temporal structure. Pezaris et al. (1997) looked for oscillations using autocorrelation functions but averaged over different behavioral conditions and reported their absence. Pezaris et al. (1999) showed structure in auto- and cross-covariation between spike trains without establishing statistical significance. LFP activity has not been previously studied in parietal cortex during working memory. We use multitaper spectral analysis techniques, which were important to our obtaining the present results, to investigate temporal structure in spiking and LFP activity. We find significant structure in the spectrum of spiking and LFP activity and the coherency between them. These results are evidence for dynamic memory fields specific to both the direction and time of a planned movement and suggest temporal structure in neuronal activity may reflect neural processing and could be used in the control of a neural prosthesis.

3.2 Materials and methods

3.2.1 Animal preparation

Recordings were made from two adult male Rhesus monkeys (*Maccaca mulatta*). Each animal was fitted with a stainless steel head post embedded in a dental acrylic head cap to fix their head position, a scleral search coil to record eye position and a stainless steel recording chamber over a craniotomy to gain access to the cortex. All surgical procedures and animal care protocols were approved by the California Institute of Technology Institutional Animal Care and Use Committee and were in accordance with National Institutes of Health Guidelines.

3.2.2 Behavioral task

Recordings were made while animals performed a memory-saccade task (see Fig.1). Each trial of this task began with the illumination of a central fixation light, to which the animal made a saccade. As long as the fixation light was present, the animal was required to maintain fixation within a 2° circular window. After foveation was established for a background period of 1 s, a location was cued by a light flashed for 100 ms at one of 8 fixed locations evenly distributed on a circle, radius 10° . Following the location cue, the monkey had to maintain fixation for a further period of 1000 ms, at the end of which the fixation light was extinguished, and the animal was required to saccade to the remembered location. To help the animal maintain saccadic accuracy, the cued location was reilluminated for a minimum of 500 ms, often triggering a corrective saccade, and the animal was required to hold fixation at the new location while the target remained on. On completion of a successful trial, the animal was rewarded with a drop of water or juice. Cued locations were randomly interleaved to collect between 10–15 successful trials for each location in blocked fashion. Locations were cued with spots of light generated with a two-beam optical bench from incandescent sources, and rear-projected onto a screen for viewing by mirror galvanometers.

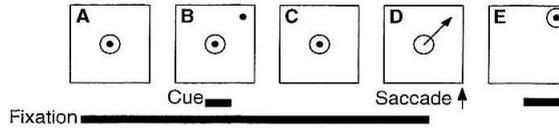


Figure 3.1: The monkey performs a memory-saccade to one of eight saccade directions. a) The trial begins with the illumination of a fixation light at the center of the screen. The monkey saccades to the fixation light and maintains fixation for one second, which determines the baseline period. b) A location is then cued at one of eight points by a flash of light for 100ms and extinguished. c) The monkey must maintain fixation for a further second at which point d) the fixation light is extinguished and the monkey performs a saccade to the remembered location cue. When the saccade is completed and the monkey’s eye position is within 2° of the target, e) it reilluminates and the monkey is rewarded with a drop of juice. A small corrective saccade follows the reillumination of the location cue. There is a short (2–3 s) intertrial interval before the fixation light reilluminates signalling the start of a new trial.

3.2.3 Electrophysiological recordings

Electrical activity was recorded from the behaving monkey using single tetrodes (Reece and O’Keefe, 1989) adapted for use in the awake monkey preparation (Pezaris et al., 1997). This is extensively described by Pezaris (2000), and is briefly summarized here. Tetrodes, made from $12\mu\text{m}$ tungsten wire (California Fine Wire, CA), were placed in a fine guide tube and positioned using a standard hydraulic microdrive (Fred Haer Corp, Brunswick, ME). Neural signals were amplified by a custom four-channel headstage amplifier (gain of 100) feeding a custom four-channel variable-gain preamplifier (gain of 1–5000 nominally set to 200 (Pezaris, 2000)) and anti-alias filters (9-pole elliptical low-pass, $f_c = 10$ kHz, Tucker-Davis Technologies (TDT), Gainesville, FL) before being digitized with a four-channel instrumentation-grade 16-bit analog to digital converter ($f_s = 20$ kHz, also TDT). Digital data were then streamed to disk and written to CD-ROM. The polarity of the signal was reversed to give positive-going spike activity. Continuous extracellular traces were processed off-line to extract and classify spike events and calculate the LFP. Figure 3.2 presents the extracellular potential from a channel of the tetrode. Panel a) presents activity during one trial and panel b) presents activity during the memory period on an expanded time base.

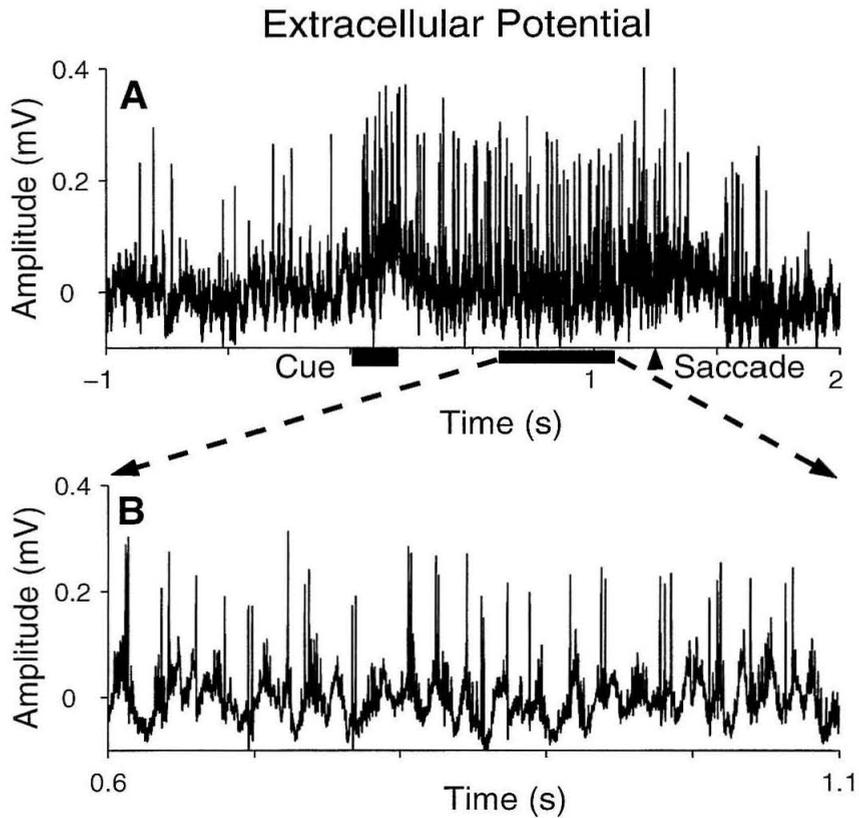


Figure 3.2: a) One channel of broadband (1 Hz –10 kHz) tetrode data sampled at 20 kHz for a memory-saccade trial during a saccade to the preferred direction. The polarity of the potential is reversed. b) The data in panel a) viewed on an expanded time base during working memory from 0.6–1.1 s.

3.2.4 Spike sorting

Spike activity was extracted from the digitized recordings by an automated procedure which identified and sorted spike waveforms into clusters, each presumed to arise from a single cell. The algorithm has been developed and described by Sahani et al. (1998) and in detail by Sahani (1999) and is summarized here.

Prior to spike sorting, the recordings were bandpass filtered (0.6–10 kHz). A statistical model describing the distribution of spike shapes was then fitted to waveforms extracted from 30 s of data, as described below. This model was used to classify spike events in the rest of the recording.

First, candidate spike times were identified by comparing the signal to a threshold of three times the root-mean-square (RMS) signal value on each channel. Spikes were accepted when the trace i) crossed and remained above the threshold for at least 0.1 ms on at least one channel; ii) crossed the threshold on the other channels either within 0.1 ms of this time, or else not within 1 ms; iii) did not remain above threshold for longer than 1 ms; and iv) did not cross the threshold again within 1 ms. These constraints reduced the number of overlapped spike events in order to reduce bias in estimating the spike waveform model.

A 2.4 ms segment (48 samples per channel) of data was then extracted from all four channels, centered on each identified spike time. A (two times RMS) thresholded center of mass was calculated for each spike waveform, and the segment resampled by interpolation to yield 24 samples per channel (1.2 ms), with the center of mass falling one-quarter of the way into the waveform. The precision of this center of mass alignment was 4–8 times the original sampling frequency. The different channels were then concatenated to yield 96-dimensional event vectors.

The background noise covariance expected in these event vectors was estimated using 1.2 ms segments extracted from the recording at times when no threshold-crossing was seen. The event vectors were then transformed to lie in a space where this noise covariance was whitened, and a mixture of a single Gaussian and a uniform density fit to them. The principal eigenvectors of the covariance of the Gaussian provided robust estimates of the principal

components of the data in the noise-whitened space. Low-dimensional event vectors were obtained by projecting each transformed event onto the four leading eigenvectors of the Gaussian covariance.

Events were clustered by fitting a mixture of Gaussian distributions to the low-dimensional event vectors. The covariance of each Gaussian was fixed at the identity matrix (since the background noise is white in the transformed event space). A uniform component was introduced in the mixture to reduce the effect of outliers. The fit was performed using the Relaxation Expectation-Maximization algorithm (Sahani, 1999; see also Ueda and Nakano 1994) which helped to avoid problems of local minima. The number of Gaussian components was determined by cascading model selection (Sahani, 1999).

Each Gaussian component in the mixture was taken to represent the distribution of spikes expected from a single cell, and events were assigned to cells according to a maximum *a posteriori* rule. The autocorrelogram of spikes assigned to each cluster was checked to ensure that no violations of the biophysical refractory period (<1 ms) were seen. To ensure that the clustering was robustly determined, the segment of data used to fit the model was varied and only models that were consistent for all segments were included in the database.

3.2.5 Data analysis

Spikes were binned at 1 ms time resolution to give spike trains. LFP time series were calculated from the extracellular recording on one tetrode channel by low-pass filtering the signal at 250 Hz (Thomson, 1994). LFP records were inspected for evidence of artifactual power at 60 Hz and its harmonics by the presence of sharp spectral lines at those frequencies. An artifact was present in some recordings during periods of low signal strength but had little power and its influence could not be detected in subsequent data analysis. When correlating spiking and LFP activity we were concerned about spikes leaking through the filter into the LFP. We took two steps to reduce those artifacts. First, we estimated the LFP from the channel of the tetrode that had the smallest spike amplitude. Second, we subtracted a mean spike waveform from the extracellular recording at each spike time before filtering.

MEAN RESPONSES: The mean response of spike activity, the peri-stimulus time histogram (PSTH), was calculated by counting the number of spikes per 1 ms bin and averaging across trials for each saccade direction aligned to the location cue onset, followed by smoothing with a Gaussian kernel, ($\sigma = 10$ ms). The mean LFP response was calculated by averaging the amplitude of the LFP across trials for each saccade direction aligned to the location cue onset. This was then smoothed with a Gaussian kernel, ($\sigma = 10$ ms). The mean response was also calculated by aligning to the time of the saccade before averaging (Barash et al., 1991) with little change in the results.

SPECTRAL ANALYSIS:

Three periods during the trial were investigated: *baseline*, *working memory* and *peri-saccadic*. The baseline period extended 750–250 ms before the onset of the location cue. The working memory period extended 450–950 ms following the offset of the location cue. The peri-saccadic period extended 250 ms either side of the saccade. Baseline activity was estimated by pooling activity from all successful trials. Working memory and peri-saccadic activity were estimated by pooling activity from successful trials according to saccade direction.

Two saccade directions of interest were defined for spiking and LFP activity, *preferred* and *anti-preferred* (opposite to preferred). The preferred direction of spike activity was the direction that elicited the maximum firing rate during the period from the offset of the location cue to the offset of the fixation light. The preferred direction of the LFP was defined as the direction with the greatest gamma band (25–90 Hz) power during the same period. In this study, multiple cells recorded at the same site had the same preferred direction and the preferred direction of spiking and LFP activity were in alignment in all recordings. Activity from different locations was aligned to the preferred direction before estimating population quantities.

Time-frequency representations of the activity were calculated on a 500 ms window that was stepped by 50 ms between estimates through the trial with the time index aligned to the center of the analysis window. As a control, window size was varied from 200–750 ms

with no significant change in the results. Time-frequency representations were estimated by averaging trials aligned in time to the location cue.

SPIKE SPECTRUM: The spike spectrum was estimated on a 500 ms window using 9 Slepian data tapers with time-bandwidth product $2p = 10$ giving a frequency resolution of ± 10 Hz. The result was then smoothed with a lag window with the same bandwidth for visualization purposes (Percival and Walden, 1993). The spike spectrum was estimated for each cell giving the *single cell* estimate. Spike activity from each cell was aligned according to preferred direction and averaged for each saccade direction across all cells in each monkey to give the *population average*. We normalized the spectrum by dividing by the mean firing rate to give the *rate-normalized spectrum* when comparing across cells or periods of the task. Spike spectrum significance levels were computed from estimates of the variance using a jackknife over tapers and trials. (Thomson and Chave, 1991; Efron and Tibshirani, 1993) and were used to test the hypothesis that the spectrum was non-uniform.

LFP SPECTRUM: The LFP spectrum was estimated on a 500 ms window using 5 Slepian data tapers with time-bandwidth product $2p = 5$ giving a frequency resolution of ± 5 Hz. The result was then smoothed with a lag window of the same bandwidth for visualization. The spectrum of LFP activity was estimated for each recording giving the *single site* estimate. LFP activity from each site was aligned according to preferred direction and averaged for each saccade direction across all sites in each monkey to give the *population average*. LFP spectrum significance levels were estimated with multitaper methods using the jackknife, as for the spike spectrum, and used to test the hypothesis that the LFP spectrum was different in one period of the trial compared with another.

SPIKE-FIELD COHERENCY: Spike-field coherency was estimated on a 500 ms window using 9 Slepian data tapers with time-bandwidth product $2p = 10$ giving a frequency resolution of ± 10 Hz. Identified spike activity was suppressed in the extracellular potential before estimating the LFP by subtracting a 2 ms mean spike waveform aligned to the event time. The spike sorting stability criterion described above insured the effectiveness of this suppression.

The coherency was estimated for each cell giving the *single cell* estimate. Spike-field coherency estimates were aligned according to preferred direction and the complex coherency values averaged to give the *population coherency*. Significance levels were estimated with multitaper methods using the jackknife, analogously to the spectrum.

SPIKE-TRIGGERED AVERAGE (STA) POTENTIAL: The relation between spike activity and the extracellular potential was assessed with the use of a STA of the raw data during baseline and working memory prior to a saccade in the preferred direction for that cell. Segments of the raw voltage trace 200 ms long centered on the spike were extracted. Spike waveforms were suppressed by subtracting a 2 ms mean spike shape waveform. The traces were then averaged to give the STA potential. The STA potential provided a measure of association between spike activity and the extracellular potential in the time domain. Since any residual spike energy would be visible as a sharp (< 1 ms) fluctuation in the spike-triggered potential at zero time, this allowed an assessment of the spike suppression procedure. 95% confidence intervals were calculated by estimating the standard error of the mean and were used to test the hypothesis that the STA potential was non-zero.

PHASE HISTOGRAM: The LFP was filtered into the gamma frequency band (25–90 Hz) using a 100 ms digital filter constructed from 6 Slepian functions ($NW = 3.5$). This gave a complex-valued series. The phase of this signal was sampled at spike times during the baseline and working memory and pooled from the preferred direction during working memory, and from all trial conditions during baseline then, finally, pooled across all cells in the study. These phases are the frequency domain version of the STA potential. The histogram of phases between baseline and working memory provided a control for the coherency as estimation of the histograms was not limited to spikes from a particular cell. A one-sample Kolmogorov-Smirnov (K-S) test was used to determine significance level for deviations of the phase histogram during working memory or baseline from a uniform distribution. A two-sample K-S test was used to determine the significance level for deviations of the phase histogram during working memory from baseline (Rao, 1965).

3.3 Results

The database for this study contained 16 cells recorded at 16 sites in one monkey and activity from 24 cells recorded at 18 sites in another monkey. The LFP was extracted from all these recordings (16 sites in one monkey, 18 sites in the second). Since we were primarily interested in memory period activity, a subjective evaluation was made during data collection to make recordings when tuned memory activity was present in multiunit activity. Subsequently, during off-line analysis, recordings were further selected for containing at least one clearly and stably isolated cell, using criteria described in the Methods section above.

3.3.1 Mean spike and LFP responses

Mean spike activity in area LIP during memory-guided saccades has been characterized in previous reports (Gnadt and Andersen, 1988; Barash et al., 1991). We found 28 of 40 neurons (70%) had significant memory period activity ($p < 0.05$).

The upper panel of Figure 3.3a presents line plots of the spike rate in the preferred and anti-preferred directions for a typical cell with memory activity. The PSTH is shown across trial conditions in a 2D plot in the lower panel. The upper panel of part b) shows the mean LFP response for the same site in line plots for the preferred and anti-preferred directions. This was calculated by averaging the amplitude of the LFP across trials to the same direction. The response across trial directions is shown in a 2D plot in the lower panel. Increases in the mean LFP are spatially tuned during and just after the location cue like those in the firing rate but spatially tuned memory activity is not. Broad tuning of the mean LFP is present across saccade directions peri- and post-saccadically which may be related to the updating of movement fields following the saccade. There is also a suppression in the mean LFP response following the saccade that is spatially tuned. These features are seen in recordings from all sites in area LIP in both monkeys and are evidence for behaviorally-locked responses in the LFP at low frequency.

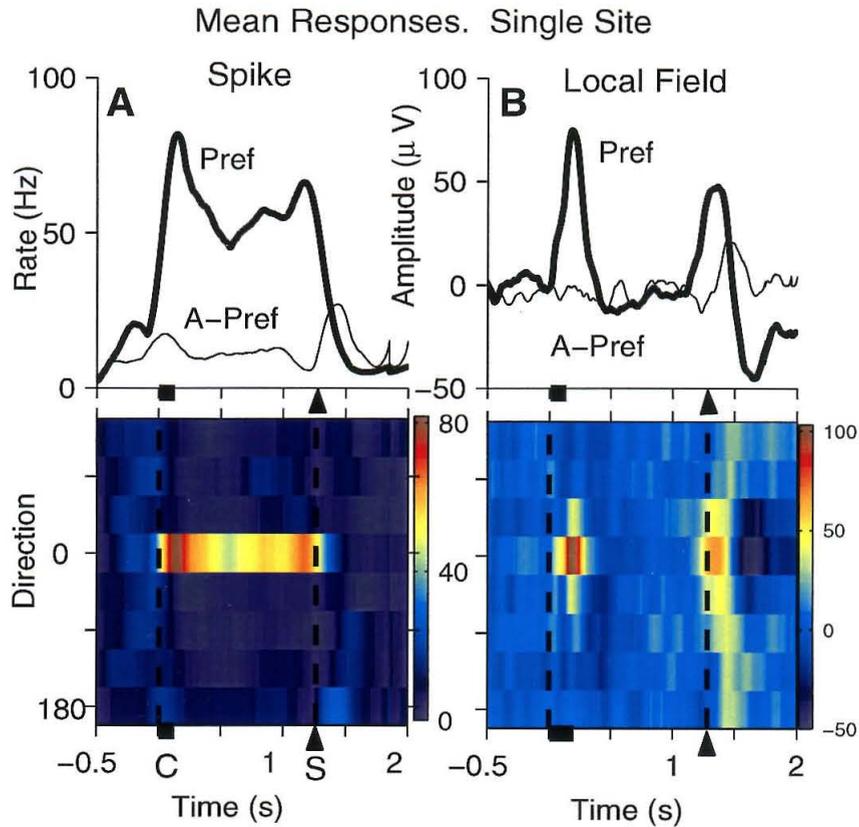


Figure 3.3: Line and 2D plots showing the mean response of spike and LFP activity averaged across saccades in each direction. a) Upper panel: Line plots of spike mean response in preferred (solid) and anti-preferred (dashed) directions. Lower panel: 2D plot of the PSTH averaged across trials aligned to the initial target onset as a function of cue direction. b) Upper panel: Line plot of the LFP mean response in the preferred (solid) and anti-preferred (dashed) directions. Lower panel: 2D plot of the LFP mean response averaged across trials aligned to the initial target onset as a function of cue direction. Time is on the x axis and for 2D plots, direction is on the y axis aligned to the preferred direction at 0° .

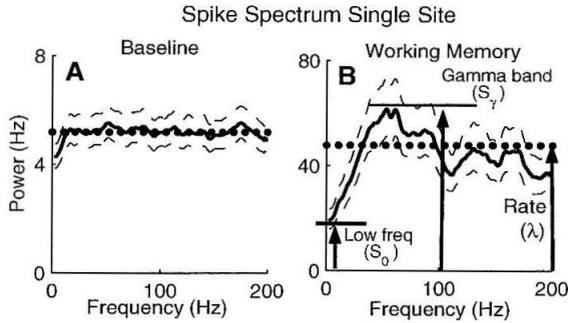


Figure 3.4: Line plots showing the spectrum of spike activity for a single cell averaged during baseline and working memory. a) Line plot of the spectrum during the baseline period (solid) with 95% error bars (dashed) estimated with a jackknife across tapers and trials. The high-frequency limit is shown by the dotted line. b) Line plot of the spectrum during working memory. Three features are indicated: 1) The value of the spectrum at low frequencies S_0 ; 2) the value at the peak in the gamma band S_γ and 3) the mean firing rate, λ . These measures are combined to give a measure of the spectral shape in the gamma band, δS (see main text).

3.3.2 Temporal structure in spike activity

Increases in the mean firing rate of cells during working memory tasks have been cited as evidence for memory fields that predict saccades to a remembered location in space (Gnadt and Andersen, 1988; Funahashi et al., 1989). Temporal structure in spike activity may also predict saccades to remembered spatial locations which would be evidence for memory fields of temporal structure, or dynamic memory fields. Here, we define dynamic memory fields in spike activity by changes in the shape of the spectrum during working memory. In contrast, memory fields described by the mean firing rate are defined by changes in the high-frequency limit, or level, of the spectrum, not the shape. Other measures of temporal structure constructed from an interval representation will be discussed below.

Figure 3.4 presents spectral analysis of spike activity from a typical cell with memory activity. Panel a) shows the results during the baseline. The level of the spectrum is equal to its mean firing rate and the spectrum of a Poisson process with the same mean rate would fall on this line. The measured activity is not significantly different from Poisson except at low (< 10 Hz) frequencies where it is suppressed.

Figure 3.4b shows the results of the analysis of activity during working memory from the same cell. The spectrum during working memory is significantly different from a Poisson process. There is a significant peak in the spectrum in the gamma band at 50 Hz. The spectrum during working memory is also strongly suppressed at low frequency (0–10 Hz). These two features indicate the cell fires spikes at regular intervals during working memory. The figure is annotated with the peak of the spectrum in the gamma band (25–90 Hz), S_γ , the value of the spectrum at 0 Hz, S_0 , and the mean firing rate λ .

We combine these features into a spectral shape parameter, δS , that is used to quantify the degree of temporal structure in the spike spectrum across a population of recordings.

$$\delta S = \frac{S_\gamma - S_0}{\lambda}$$

This measure is related to the derivative of the spectrum and quantifies deviations of the spectrum from a Poisson process that are local in frequency. Importantly, since δS is explicitly normalized by the level of the spectrum it is a dimensionless number and is not biased by the mean firing rate.

Figure 3.5a presents results during baseline and panel b) presents results during working memory. The top panel shows the population average spectrum in each case. 95% confidence intervals shown by dashed lines indicate that deviations from the Poisson level are significant during working memory across the population and not during baseline. The lower panels of Figure 3.5 present the distribution of the spectral shape parameter pooled across all cells in the database.

The central vertical panel shows the distribution of spectral shape for baseline and working memory activity across the data set. These distributions were unimodal, but somewhat skew so we used the median as a measure of central tendency, displayed as a notch in the box plot for each condition. Robust 95% confidence intervals for this estimate are shown by the extent of the notch about the median. Similar box plot displays present population data in later figures.

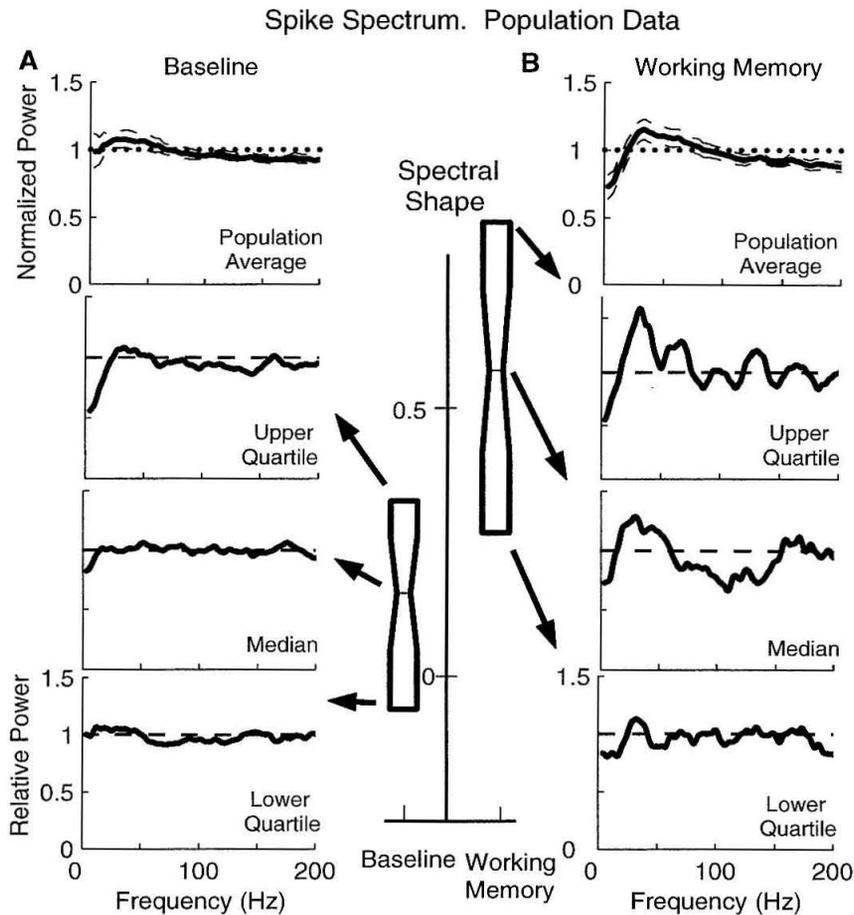


Figure 3.5: A comparison of spectral structure in spike activity during baseline and working memory across a population of recordings. Central panel: Box plots showing the spectral shape for spike activity, δS , during working memory in the preferred direction compared to the baseline, pooled across a population of recordings from both monkeys. δS is unimodal and the distributions are skew so the median is used as a measure of central tendency, shown at center of the box plot for each condition. Robust 95% confidence intervals for this estimate are shown by the extent of the notch. *Side panels:* Line plots showing the rate-normalized spectrum out to 200 Hz during baseline (left) and working memory (right). Individual plots of spectra in the upper, median and lower quartiles are shown, as well as the population average spectrum for baseline and working memory activity. Arrows indicate the approximate location of each spectrum in the distribution. The high-frequency limit is shown by dotted, horizontal line at 1.

The median spectral shape during the baseline is small indicating that baseline activity contains little temporal structure. In contrast, the distribution of spectral shape during working memory is significantly larger ($p < 0.05$). This is consistent with the population average spectrum and shows that working memory activity in area LIP contains significant temporal structure not present during the baseline.

The individual plots of the spectrum during baseline reveal little temporal structure except for a small amount of suppression at low frequencies (0–10 Hz). In contrast, those during working memory reveal significant temporal structure with a peak in the gamma frequency band (25–90 Hz) and sharp suppression at low frequencies (0–10 Hz) across the distribution.

Analysis of temporal structure in non-preferred directions suggests dynamic memory fields are aligned to the same spatial location as fields defined by the mean firing rate. Under the present experimental design, equal numbers of trials are performed for each cue direction, resulting in data from the non-preferred directions having fewer spikes available for analysis. Since resolving temporal structure in spiking requires more spikes than estimating the mean firing rate, it is difficult to estimate how the spectrum changes with the rate in these data. Figure 3.9, below, presents the distribution of spectral shape for each direction across the population and demonstrates this problem.

We estimated the spectrogram of spike activity during trials to the preferred direction to see how temporal structure is organized during the trial. Figure 3.6a presents the results for a typical cell and panel b) after averaging across the population of cells. The spectrogram, shown in the upper panels, is dominated by changes in the rate. To compare temporal structure between time intervals with different firing rates, we present an estimate of the rate-normalized spectrogram in the lower panels. The rate-normalized spectrogram suggests the temporal structure during working memory is not a visual response to the location cue since it begins after the offset of the location cue and extends through the saccade.

Additional analysis not presented here show that the shape of the spectrum given by the frequency at the gamma band peak is relatively constant across the population of cells in both monkeys and is not correlated with the the mean firing rate. This is evidence the

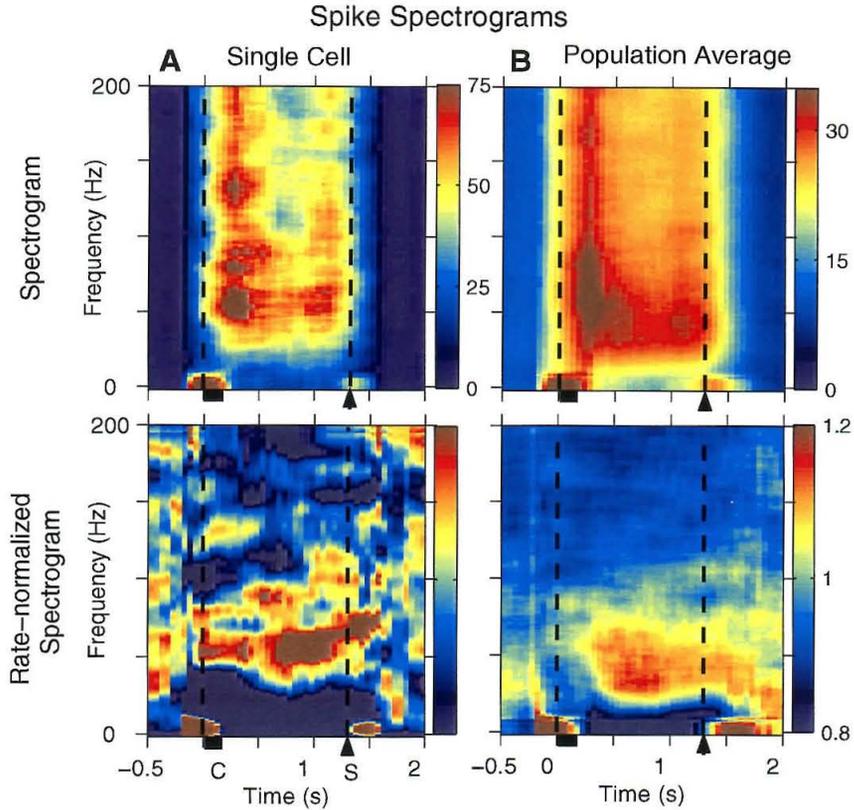


Figure 3.6: 2D plots showing the spectrum of spike activity in the preferred direction with and without rate normalization for a single cell and a population average across both monkeys. a) Upper panel: 2D plot of the spectrogram for a single cell. Lower panel: 2D plot of the spectrogram for a single cell normalized by the rate. b) Upper panel: 2D plot of the population average spectrogram. Lower panel: 2D plot of the population average spectrogram normalized by the rate. Time is on the x axis, frequency on the y axis and power is color-coded on a linear scale.

spike activity is not well modelled by a simple integrate-and-fire model with uncorrelated inputs. The mean firing rate of the cell is, however, correlated with the strength of the spectral suppression and its width in frequency.

These data contradict the commonly held assumption that spike arrival times are independent and are evidence for the presence of significant temporal structure that is modulated by the behavioral state of the monkey. During the baseline, the probability of the cell firing can be modelled as a Poisson process. In contrast, the activity during working memory shows detailed temporal structure. The probability of the cell firing is suppressed for a period around each spike and then enhanced to give increased power in a broad, gamma frequency band (25–90 Hz).

3.3.3 Temporal structure in LFP activity

LFP activity is generated by extracellular currents that are thought to be related to synaptic activity in a local population of cells (Mitzdorf, 1985). Previous studies have reported that temporal structure is present in LFP activity in cats (Gray and Singer, 1989) and macaques (Eckhorn et al., 1993; Kreiter and Singer, 1996). If the temporal structure we observe in spiking during working memory is a result of synchronized synaptic activity it may also be present in the LFP.

Figure 3.7 presents the LFP spectrum at a single site for trials to the preferred direction. Panel a) presents the results during working memory (solid) and baseline (dashed). Although most of the power in the LFP is below 30 Hz, working memory activity above 30 Hz is elevated compared to baseline and peaked in the gamma frequency band (25–90 Hz). Panel b) presents the relative power of the spectrum between working memory and baseline formed by dividing the spectrum during working memory by the spectrum during baseline. Temporal structure in LFP activity changes between working memory and baseline in a gamma band, but not at low frequencies.

The significant difference in power in the gamma band (25–90 Hz) is also present in a population average of the spectrum. Figure 3.8a presents the population average spectrum during the baseline and panel b) during working memory in the preferred direction. The

LFP Spectrum

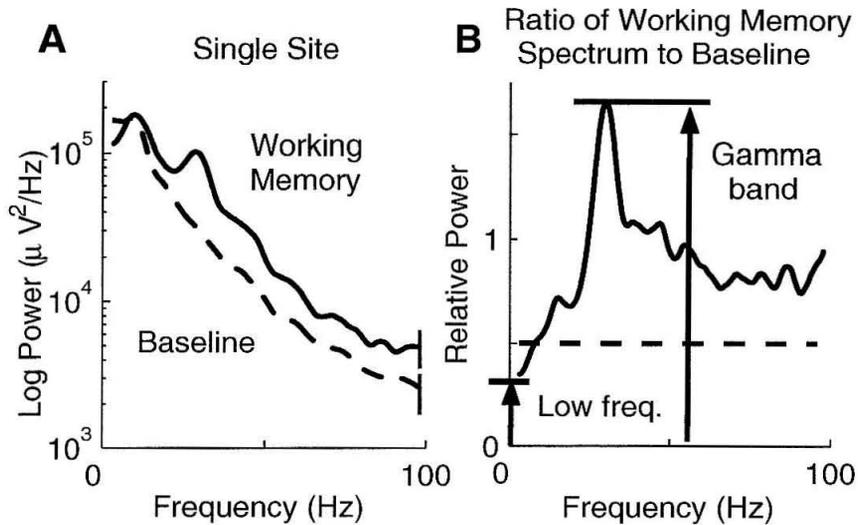


Figure 3.7: Line plots showing the spectrum of LFP activity averaged during baseline and working memory. a) Line plot of the spectrum during baseline (dashed) and working memory (solid) on a log scale at a single site. The bar indicates the extent of 95% error bars for the spectrum estimated with a jackknife across tapers and trials. b) Line plot of the ratio of LFP spectrum during working memory and baseline on a linear scale at a single site. Arrows highlight the power at low frequency (0–10 Hz) and in the gamma band (25–90 Hz). Characterization of the interval properties of spike activity from a sample cell. a) The ISI histogram during the baseline. b) The ISI histogram during working memory. c) The interval spectrum for memory activity (solid) with 95% confidence intervals (dashed). The dotted horizontal line shows the spectrum of a renewal process for comparison.

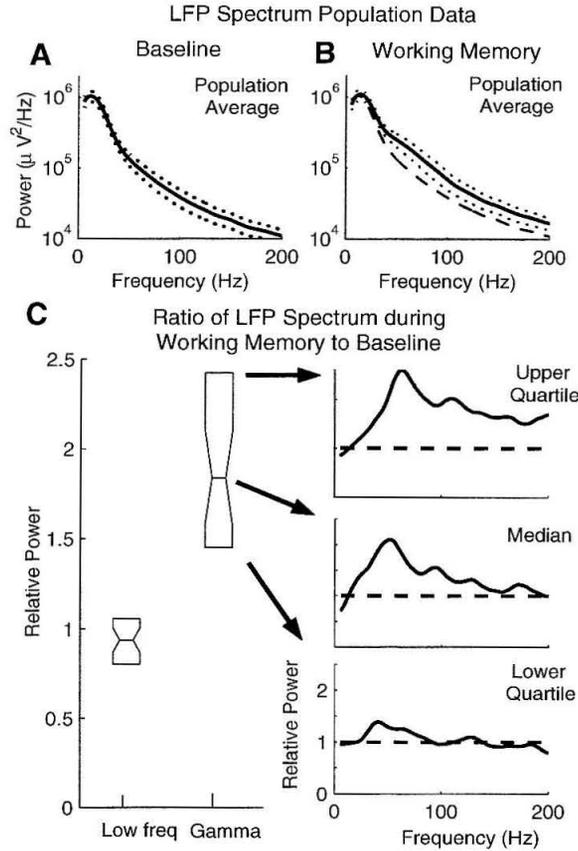


Figure 3.8: Comparison of changes in relative power between baseline and memory activity in the LFP across a population of recordings. a) LFP spectrum averaged across a population of sites recorded in two monkeys. Average spectrum during baseline with 95% confidence intervals. b) Average LFP spectrum during working memory in the preferred direction with 95% confidence intervals. c) The left panel shows box plots showing low frequency (0–10 Hz) and gamma band (25–90 Hz) power in the LFP during working memory in the preferred direction relative to baseline pooled across a population of recordings from both monkeys. The right panel shows line plots of the ratio of power during working memory to baseline against frequency out to 200 Hz. Plots of spectra in the upper, median and lower quartiles are shown and asterisks mark their approximate position on the box plot. The center of the notch marks the median and the extent of the notch gives confidence intervals for this estimate.

confidence intervals do not overlap in the gamma band (25–90 Hz) indicating the difference in power is significant across the population.

In Figure 3.8 we also compared the relative power between working memory and baseline at low frequency (0–10 Hz) and at the peak in the gamma band (25–90 Hz) across the data set. Panel c) shows box plots of the distribution in each frequency band. The distribution of relative power at low frequency (0–10 Hz) is not different from 1, while the distribution of relative power in the gamma band (25–90 Hz) is significantly greater than 1. Panel d) presents the relative power across the distribution.

We tested the hypothesis that power in the LFP in the gamma frequency band (25–90 Hz) was greater in the preferred direction during working memory compared to baseline against the null hypothesis that there was no difference. Gamma band LFP power was significantly greater ($p < 0.05$) in the preferred direction during working memory than baseline in 27 of 34 (82%) of sites recorded from.

The increase in LFP power in a gamma band during working memory is evidence for dynamic memory fields similar to those seen in spike activity, presented above. We were interested to determine how these memory fields were spatially tuned. Figure 3.9 presents tuning curves for a) the mean firing rate; b) gamma band LFP power; c) spike spectral shape and d) low frequency LFP power across a population of recordings pooled across both monkeys. The mean firing rate and LFP gamma band activity (25–90 Hz) show clear tuning. The spike spectral shape shows some tuning, but it is not well resolved since the firing rate was low in non-preferred directions. LFP activity at low frequencies (0–10 Hz) does not show significant spatial tuning. This shows spatial tuning in the LFP is local in frequency and not due to an overall increase in power.

We tested the hypothesis that gamma band LFP power was greater during working memory in the preferred direction compared to the anti-preferred direction against the null hypothesis that there was no difference. Gamma band (25–90 Hz) LFP power during working memory was significantly greater ($p < 0.05$) in the preferred direction compared to the anti-preferred direction at 28 of 34 (85%) sites recorded from.

Figure 3.10 shows spectrograms for LFP activity averaged during trials to the preferred

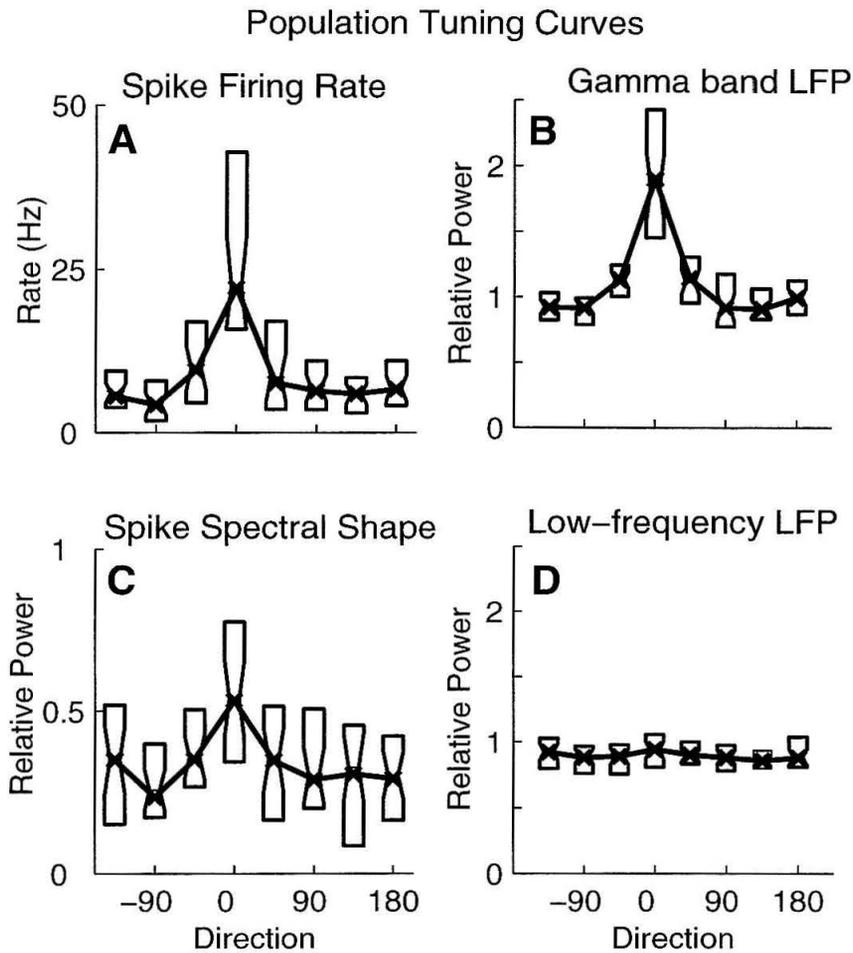


Figure 3.9: Box plots showing the tuning of spectral measures of spiking and LFP activity during the memory period relative to the baseline pooled across a population of recordings in both monkeys. a) Tuning curve with box plots showing mean firing rate across the population. b) Tuning curve with box plots showing relative power in the gamma band (25–90 Hz) across the population. c) Tuning curve with box plots of spike spectral shape, δS , across the population. d) Tuning curve with box plots showing relative power at low frequency (0–10 Hz) across the population.

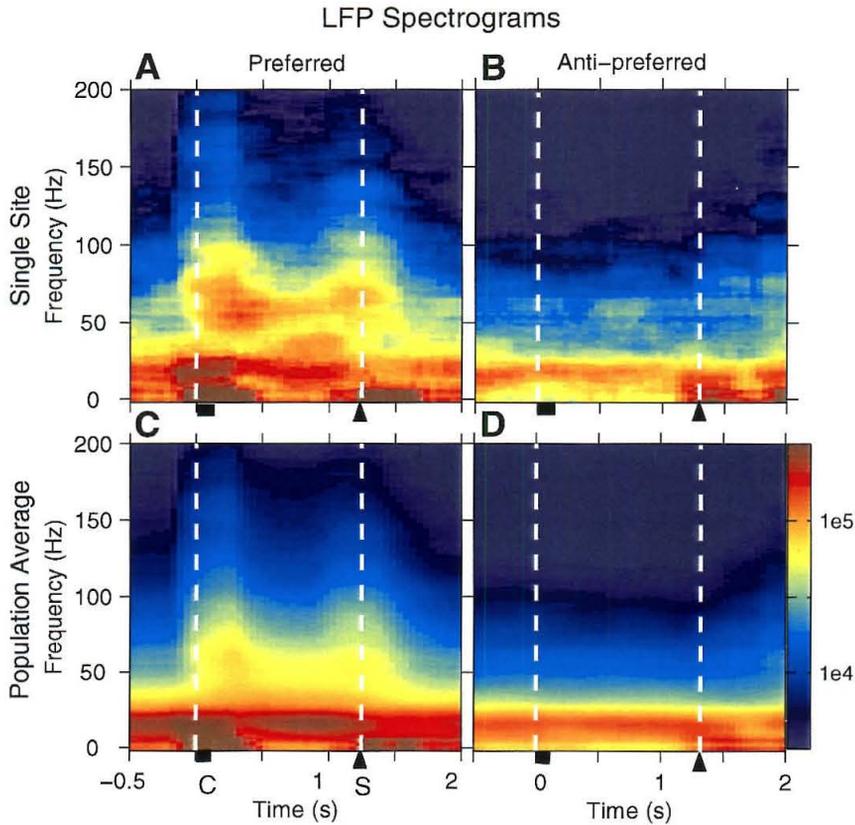


Figure 3.10: Spectrograms of LFP activity averaged across trials during saccades to either the preferred or anti-preferred direction. a) 2D plot of the spectrogram in the preferred direction at a single site. b) 2D plot of the spectrogram in the anti-preferred direction at a single site. c) 2D plot of the spectrogram in the preferred direction averaged across a population of recordings from both monkeys. d) 2D plot of the spectrogram in the anti-preferred direction averaged across a population of recordings from both monkeys. In each panel, time is on the x axis, frequency is on the y axis and power is color-coded on a log scale.

Gamma and Beta band power

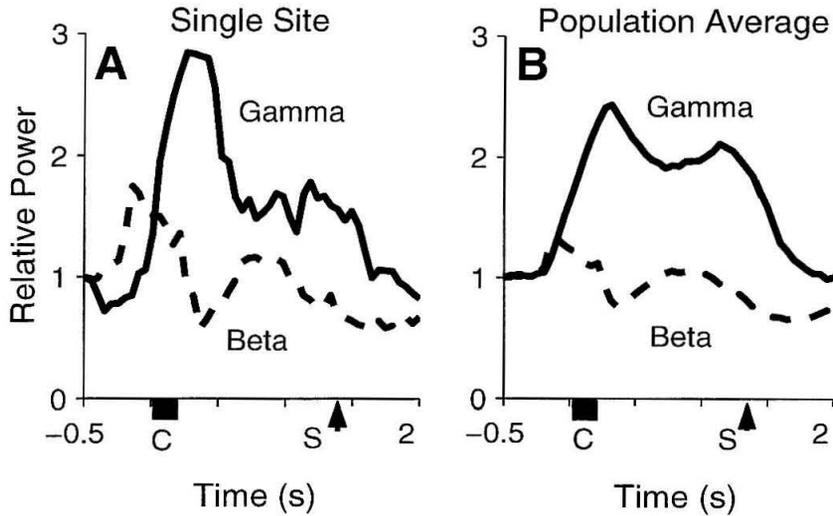


Figure 3.11: Power in the LFP in the gamma (25–90 Hz) and beta (15–25 Hz) frequency bands during trials with saccades to the preferred directions. Power in each frequency band is normalized by the level during baseline. a) Gamma band (solid) and beta band (dashed) activity from a single site. b) Population average gamma band (solid) and beta band (dashed) activity from two monkeys.

and anti-preferred directions. In the preferred direction, the increase in gamma band power is sustained during the memory period through the saccade for a single site and across the population. In addition to dynamic memory fields the LFP contains different task-related activity in the beta frequency band (15–25 Hz). This activity does not show simple spatial tuning, but is modulated during the trial.

Figure 3.11 shows LFP power in the gamma and beta frequency bands over time. Elevations in gamma band power are clearly seen and extend through the saccade, similar to the elevations of gamma band power in spiking activity. Activity in the beta band is modulated during working memory and peri-saccadically. Immediately after the location cue, beta band power is suppressed and it then increases toward the end of the memory period. Beta band activity is also suppressed peri-saccadically during saccades to all directions and this is present in activity at a single site and a population average. We tested the hypothesis

that peri-saccadic activity in the beta frequency band (15–25 Hz) is suppressed compared to working memory against the null hypothesis that there was no difference. Beta band activity was significantly suppressed ($p < 0.05$) at 34 of 34 (100%) sites recorded from.

These observations are evidence for temporal structure in the LFP in multiple frequency bands. Activity in the gamma frequency band is spatially tuned while activity in the beta frequency band is dynamically modulated during the task.

3.3.4 Coherency between spiking and LFP activity

We report temporal structure in the gamma frequency band (25–90 Hz) that is present in both spiking and LFP activity during working memory. If spike activity is being driven by synchronized synaptic activity and these currents are also measured in LFP activity the processes may be coherent. We investigated this possibility by studying the extracellular potential centered on spikes and by calculating the coherency between spiking and LFP activity during baseline and working memory. A study of spike-field coherency may also provide useful information about the nature of non-Poisson temporal structure in spike trains.

We first estimated the STA potential. The mean potential in a 200 ms window was calculated conditional on spike times during working memory and compared against baseline at a single site. Figure 3.12a shows the STA potential during the baseline period and panel b) shows it during working memory. The STA during working memory shows an oscillatory component in the average that is absent from the baseline indicating that the phase of the potential in that frequency band is coherent with spike activity. However, this structure remains just below significance across much of the window although it is clearly present. This suggests that error bars constructed in the time domain are not suitable to detect signals localized in frequency and that frequency-domain techniques may describe this activity more clearly.

We present two ways to test for correlated activity that is localized in frequency. In the first, we estimate the coherency between spiking and LFP activity. In the second, we study a frequency domain version of the STA potential. We filter the STA potential in the

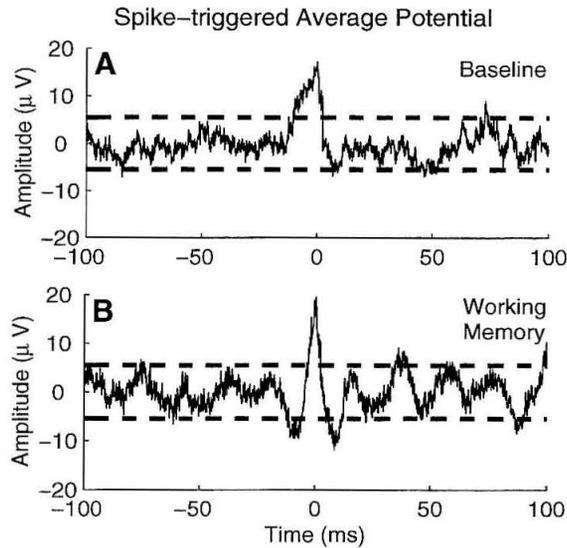


Figure 3.12: Spike-triggered average potential for activity from a single cell at a single site during baseline and working memory for saccades to the preferred direction. a) Spike-triggered average potential during the baseline. b) Spike-triggered average potential during working memory.

gamma frequency band (25–90 Hz) and compare the histogram of the phase at spike times between baseline and working memory pooled across all cells in this study.

Figure 3.13 shows the coherence between a cell and the simultaneously recorded LFP. Panel a) shows activity during baseline and panel b) activity during working memory. There is no significant spike-field coherence during the baseline whereas a sharp increase in the spike-field coherence exceeds 99% confidence intervals in the gamma frequency band during working memory.

We calculated the coherency for the 40 cells in this study during saccades to the preferred direction and compared the distribution of coherence in the gamma band between the baseline and working memory. Figure 3.14 presents the results. The central panel shows the distribution of gamma band coherence during working memory is significantly larger ($p < 0.05$) than during baseline across the population. The left panel shows no significant structure is observed during baseline; in contrast, the right panel shows there is significant coherence during working memory.

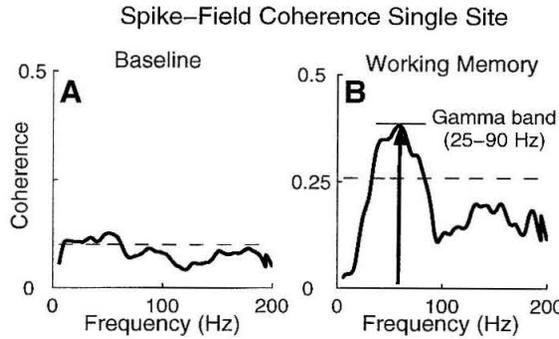


Figure 3.13: Coherency between spiking and LFP activity at a single site. a) Line plot of spike-field coherence during the baseline. b) Line plot of spike-field coherence during working memory. Dashed lines show 99% confidence intervals.

A comparison of activity between multiple cells and the LFP recorded at the same site was not possible as there were only two pairs of simultaneously recorded cells with elevations in firing rate during working memory in this study.

Figure 3.15 presents the coherence against time and frequency during saccades to the preferred direction. Panel a) shows the results at a single site and panel b) shows the population average. An increase in gamma band coherence is seen after the location cue and during the memory period. When the coherence is significant, the phase of the coherency is shown by an arrow and is approximately zero radians indicating the cell fires on or just before the peak of the LFP oscillation (according to our convention for the LFP). The phase of the coherency during working memory is relatively constant across the gamma band so spike and LFP activity do not lag each other in time. These observations are evidence for sustained phase-locking between the spiking and the LFP during working memory. The population average shows the increase in gamma band coherency between the LFP and spiking is significant ($p < 0.01$). Therefore, spiking locks to gamma band temporal structure in the LFP during working memory with the same preferred phase in all recordings from area LIP in this study.

Confidence intervals for the estimates of spike-field coherency are sensitive to the total number of spikes available, so these measures are less sensitive to coherency during baseline than working memory. We applied a test for spike-field coherency that had the same

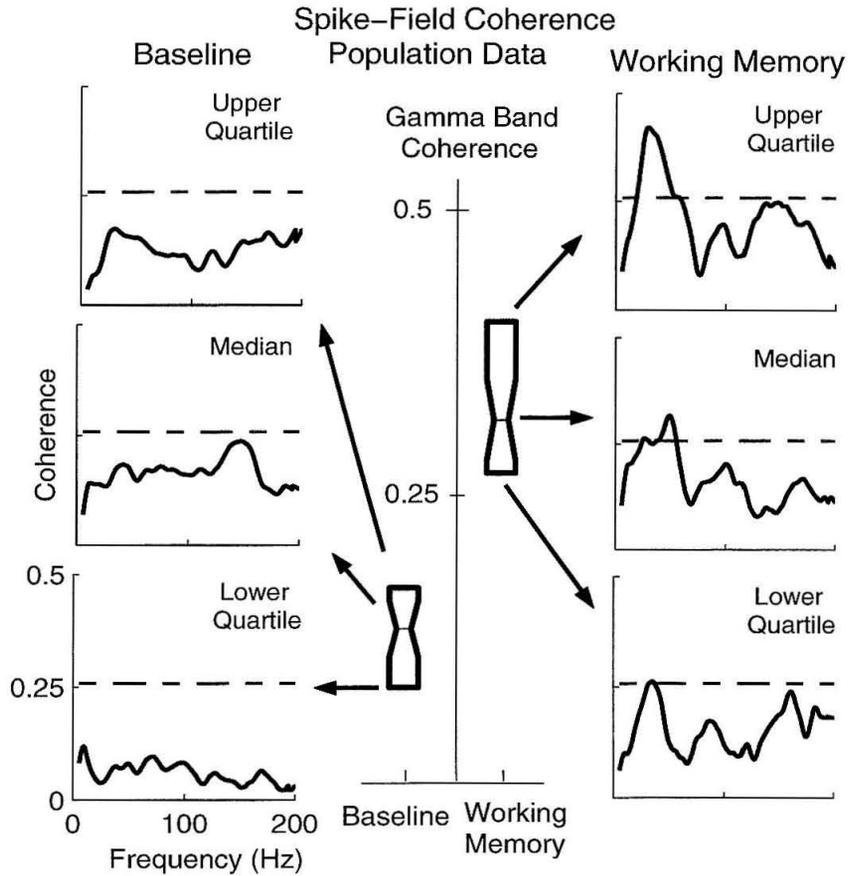


Figure 3.14: Comparison of spike-field coherence during baseline and working memory across a population of recordings in two monkeys. Central panel: Box plots showing spike-field coherence during working memory in the preferred direction compared to the baseline pooled across a population of recordings from both monkeys. Side panels: Line plots showing spike-field coherence out to 200 Hz during baseline (left) and working memory (right). Individual plots of coherence in the upper, median and lower quartiles are shown and asterisks mark their approximate position on the box plot.

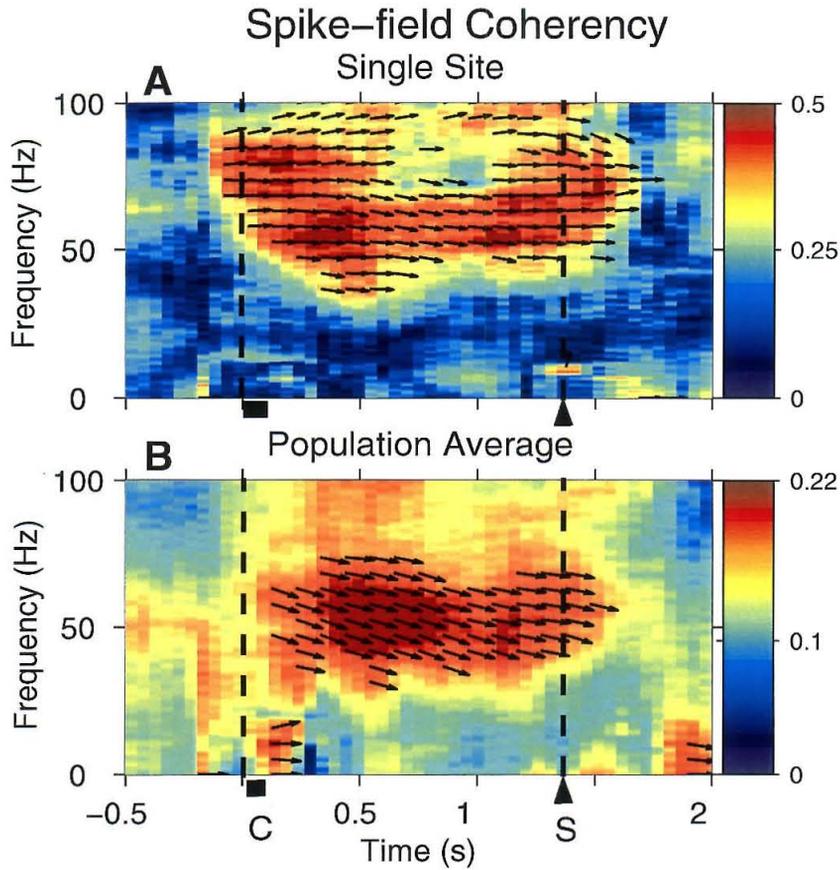


Figure 3.15: Coherency of spiking and LFP activity across time. a) Coherency at a single site in the preferred direction. b) Coherency in the population average in the preferred direction. Time is on the x axis and frequency on the y axis. The coherence is color-coded on a linear scale. Arrows denote the phase where the coherence is significant ($p < 0.01$).

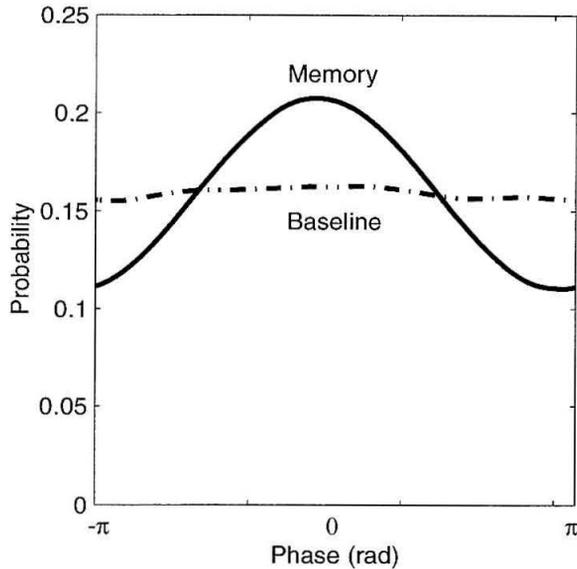


Figure 3.16: The distribution of the phase of the gamma band (25–90 Hz) activity in the LFP at spike times for all cells in the data set during saccades to their preferred direction. Working memory (solid) and baseline (dashed) period distributions are compared.

statistical power during the baseline and working memory. This is a frequency-domain version of the STA potential that could be pooled across recordings (see Section 3.2.5). Figure 3.16 shows the normalized histogram of this spike-triggered gamma band phase during working memory (solid) and the baseline (dashed). The distribution during working memory is significantly different from uniform ($p < 0.01$ K-S test: $N=6192$) as well as significantly different from baseline ($p < 0.01$ K-S test: $N=6192$) while the distribution during baseline is not significantly different from uniform ($p < 0.01$ K-S test: $N=6192$). Consequently, there is no evidence of gamma-band phase-locking between spiking and LFP activity in the gamma band during the baseline.

The presence of significant gamma band coherence between spiking and LFP activity has consequences for models of spike activity during working memory. It shows spiking during working memory cannot be modelled as a renewal process which is in agreement with results of the interval analysis presented above.

3.4 Discussion

This study investigates the relation of neuronal dynamics to working memory by examining temporal structure within and between spiking and LFP activity in area LIP of macaque parietal cortex during a memory-saccade task. We focus on the three principal findings of this work: i) spike activity contains temporal structure during working memory and not simple fixation which is evidence for dynamic memory fields; ii) LFP activity reveals spatially tuned dynamic memory fields in the gamma band (25–90 Hz) but not at low frequencies (0–10 Hz); and iii) dynamic memory fields in spiking and LFP activity are coherent in the gamma band with the same phase across LIP. This means spike activity contains temporal structure that is local in frequency. Results from the LFP suggest this activity is synchronized across large populations of neurons.

3.4.1 Understanding the temporal structure in spike activity

Our findings show spike activity during working memory is not well modelled by a Poisson and contains temporal structure local in frequency. This contradicts commonly held assumptions about the independence of spiking activity (Shadlen and Newsome, 1994). Our analysis of the spike spectrum shows that during working memory there is temporal structure due to a combination of spectral suppression and a gamma band spectral peak. The significant spike-field coherency confirms this and shows that during periods of elevated firing the temporal correlations are local in frequency. The role of temporal structure in neural coding has been controversial (Roskies, 1999). Our results are clear evidence that temporal structure in spiking activity in parietal cortex is modulated during behavior and organized in dynamic memory fields that may reflect neural processing.

3.4.2 Temporal structure in the LFP

Our results indicate gamma band LFP activity is also organized in dynamic memory fields during working memory. The tuning of temporal structure in spike activity was not well resolved in some cue directions but by analyzing the LFP we were able to study temporal

structure in all directions. This illustrates the improved statistical power of LFP activity compared with spike activity when the firing rate is low. The presence of dynamic memory fields in the LFP suggests the presence of synchronized activity in large populations of neurons. This is supported by the results of spike-field coherence and suggests the LFP reflects the columnar organization of area LIP (Blatt et al., 1990).

Two other aspects of LFP activity related to the task were also interesting. Firstly, the tuning of the LFP mean response broadens peri- and post-saccadically. The significance of this is not known, but it may be related to signals to update the eye fields in area LIP following the saccade (Gnadt and Andersen, 1988; Duhamel et al., 1992; Batista et al., 1999; Snyder et al., 1999). Secondly, the LFP has additional temporal structure in the beta frequency band that could be related to movement execution and preparatory aspects of the task. These results indicate temporal structure in the LFP in parietal cortex contains a variety of specific information about the direction and time of planned movements.

Parietal cortex is implicated in higher cognitive function and spiking activity in various parietal areas codes for movement plans (Andersen, 1995; Bracewell et al., 1996; Snyder et al., 1997). There is interest in decoding cortical signals for a neural prosthesis but acquiring spiking activity with chronically implanted electrodes remains a challenge (Shenoy et al., 1999). Since the LFP is easier to acquire than spiking, our findings suggest that decoding the temporal structure in LFP activity could accelerate the development of this application.

3.4.3 Coherent gamma band activity during working memory

We find that during working memory the coherency between spiking and LFP activity is significant in the gamma band. Spikes exhibit phase locking to the LFP and preferentially fire at the peak of a broad LFP oscillation throughout the memory period, and this temporal structure is not present during simple fixation. While we have observed coherent activity during working memory, it is important to note such activity need not be specific to working memory. Area LIP is activated by a number of other oculomotor tasks that require spatial cognition (Colby and Goldberg, 1999) and coherent gamma band activity could be present

during those tasks.

Our finding of correlations in a broad gamma frequency band modulated by behavior complements recent reports of gamma band activity during attention in the monkey (Fries et al., 2001) and working memory in human EEG (Tallon-Baudry et al., 1998, 1999). These and other studies suggest that, in addition to providing important statistical advantages, investigating correlations between spikes and LFP activity may help bridge the gap between potentially related results showing modulated temporal structure in EEG activity in humans and single unit activity in non-human primates (Tallon-Baudry and Bertrand, 1999).

Chapter 4 Interspike interval correlations in cortical activity: Data analysis and modeling

4.1 Introduction

In Chapter 3 we report on temporal structure in spiking activity and the LFP recorded in area LIP of a Macaque performing a memory task. Using spectral analysis we find spatially tuned activity in the gamma band (25–90 Hz) that is coherent between spiking and the LFP. Gamma band activity has also been reported in Macaque visual cortex in areas V1 (Friedman-Hill et al., 2000), V4 (Fries et al., 2001) and MT (Bair et al., 1994; Kreiter and Singer, 1996), as well as recently, in the parietal reach region (PRR) (Buneo et al.). Although it is widespread, the nature of the activity is debated. Some argue the gamma band spectral peak results from an oscillatory process (Singer and Gray, 1995) while others argue it is simply due to bursting (the firing of action potentials in quick succession) and not oscillatory (Bair et al., 1994).

Importantly, the spectrum or auto-correlation function is often used to characterize temporal structure in spiking. As discussed earlier (see Chapter 1.4) these estimate correlations in the counting process. In some cases, the interspike interval (ISI) histogram is also estimated, but correlations in the interval process have not been characterized in the behaving monkey before. As we will show, certain ISI correlations are expected if spiking activity participates in oscillatory processes.

In this chapter we present spiking activity from three areas that contain gamma band activity, LIP, PRR, and MT, during three different tasks, the memory-saccade, the memory-reach and a visual motion discrimination task, respectively. We extend previous analysis of these data by estimating ISI correlations using the interval spectrum. This sheds new light on the nature of gamma band activity and we discuss the implications for stochastic

process models.

4.2 Methods

4.2.1 Behavioral paradigms

Memory-saccade task

The memory-saccade task is presented in Chapter 3.2.2. In brief, animals were head-fixed in front of a vertical screen. Visual stimuli were generated by incandescent light sources. Trials began with the illumination of a central location which the animal fixated. A spatial location was cued with a flash (duration 100 ms) at one of equal angles on a circle radius 10° . After a delay period of 1000 ms, the central fixation light was extinguished and the animal made a saccade to the remembered location of the cue in complete darkness. The animal was then rewarded with a drop of juice.

Memory-reach task

The memory-reach task is the analog of the memory-saccade task, but for reaching. Full details of the task are presented in Snyder et al. (1997). Animals were head-fixed in front of a vertical board of touch sensitive buttons containing red and green LEDs. The red LED instructed where to direct their gaze while the green LED instructed animals where to place their hand. Trials began with the illumination of both the red and green LED directly in front of the animals. After the animals looked and reached to the target, a spatial location was cued for a reach by flashing (duration 300 ms) the green LED at one of eight peripheral locations. After a delay period of 800–1000 ms, the central LEDs were extinguished and the animal made a reach to the remembered location of the cue while maintaining fixation in complete darkness. The animal was then rewarded with a drop of juice.

Visual motion discrimination task

Full details of the task are presented in Britten et al. (1992). In brief, spiking activity was recorded during the performance of a visual motion discrimination task. The motion stimulus was created by coherently moving a fraction of dots in a random dot display in a particular direction. The fraction of dots, or motion coherence, ranged from 0% to 99.9%. Displays with motion coherence between these limits induced a motion percept with varying strength. Animals were trained to discriminate opposite directions of motion during a 2s period and indicate their judgement of motion by making a saccade to a target LED following the stimulus offset. They were rewarded for each correct discrimination with a drop of juice.

4.2.2 Experimental recordings

Area LIP

Single cell and LFP recordings were obtained simultaneously from sites in area LIP from two monkeys using custom tetrodes. Spike trains were binned at 1 ms resolution. Results are presented from the analysis of 40 neurons recorded in two monkeys. Full details of the experimental recording are presented in Chapter 3.2.3.

PRR

Single cell recordings were obtained from sites in PRR from two monkeys using standard tungsten microelectrodes. The LFP was not recorded. Spike trains were binned at 1 ms resolution. Results are presented from the analysis of 98 neurons recorded in two monkeys. Full details of the experimental recording are presented in Batista et al. (1999).

Area MT

Single cell recordings were obtained from sites in area MT from four monkeys using standard tungsten microelectrodes. The LFP was not recorded. Spike trains were binned at 1 ms resolution. Results are presented from the analysis of 187 neurons selected from 253 neurons

in a neuronal database published on the internet¹. Full details of the experimental recording are presented in Britten et al. (1992).

4.2.3 Data analysis

Spike trains were analyzed by estimating the spectrum, ISI histogram and interval spectrum.

SPECTRUM: The spectrum was estimated for each motion coherence condition by averaging a rate-normalized spectrogram across time and over trials. The rate-normalized spectrogram was estimated using multitaper methods ($W = 20$ Hz) with a 250 ms moving-window shifted by 50 ms between estimates. This procedure minimized the influence of rate variations on the estimated spectrum.

ISI HISTOGRAM: The ISI histogram was estimated using spike trains binned at 1 ms then smoothed by a 10 ms Gaussian kernel.

INTERVAL SPECTRUM: The interval spectrum of spike activity was estimated by calculating the spectrum of a sequence of ten or more consecutive intervals between spikes. Three Slepian data tapers with time-bandwidth product $2p = 4$ were used. The result was then smoothed with a lag window with the same bandwidth. The interval spectrum was not computed if a sequence of ten consecutive intervals in the window was not available (i.e., the firing rate was less than 20 Hz).

The interval spectrum was estimated from the sequence of interspike intervals in a single 1s window at the end of the visual stimulus period. Multitaper methods were used ($W = 0.1$) and significance was assessed using a jackknife procedure averaged over trials and tapers. The sequence of intervals was log-transformed and normalized by the variance of the interspike interval (ISI) distribution before estimating the interval spectrum. The log transformation stabilized the variance of the ISI sequence making it more suitable for spectral analysis. When the analysis was repeated without this transformation, the results were qualitatively the same but with larger confidence intervals. This is because the

¹<http://www.cns.nyu.edu/~wyeth/data/newsome/newsome.html>

untransformed ISI distribution has a long tail.

4.3 Results

4.3.1 Stochastic process models

A variety of stochastic process models have been used to describe neuronal firing. Typically, models are matched to metrics that reflect spike count and interspike interval variability. Since we will be analyzing ISI correlations in experimental data, we focus on models described by their interval properties and their spectra.

Elementary models

The sequence of time intervals between consecutive spikes, or interspike intervals (ISIs), forms a representation of spike activity, the interval representation, that is complementary to the sequence of spike arrival times, or counting representation, that we have considered so far. Activity from a Poisson process, which has no temporal structure, has the following interval properties: i) the ISI distribution falls exponentially with the length of the interval and ii) the sequence of ISIs is uncorrelated. Figure 4.1 presents analysis of simulated data from a Poisson process to illustrate these properties.

A renewal process, which can contain temporal structure, generalizes the interval properties of the Poisson process by allowing the ISI distribution to take any form. (Cox and Lewis, 1966). All the temporal structure in a renewal process is represented in the distribution of times to the next spike i.e., the distribution of interspike intervals (ISIs). You can derive the spectrum from the ISI distribution. This is not true in general as the correlations between intervals make it difficult to transform joint distributions of intervals to counts and vice versa Cox and Lewis (1966). Traditionally, the renewal process has been used to model neuronal spiking activity because it provides a simple mathematical description of a cell whose membrane potential is reset after each spike, and more generally, it can model how a spike event affects the time of the next spike.

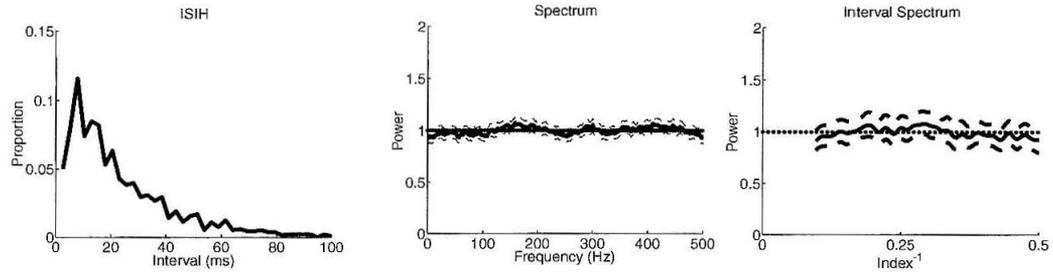


Figure 4.1: Activity from simulated Poisson process. a) The ISI histogram. b) Spectrum normalized by the rate (solid). Poisson limit (Dotted). c) The interval spectrum (solid) with 95% confidence intervals (dashed). The dotted horizontal line shows the spectrum of a renewal process for comparison.

Two kinds of temporal structure widely observed in cortical spiking activity are refractoriness (in addition to the biophysical refractory period) and bursting. The biophysical refractory period, in which the cell is unable to fire an action potential for ~ 1 ms after each spike is well-established, but its effect is subtle at firing rates less than 100 Hz. Refractoriness describe the tendency of a cell not to fire for up to 20 ms following each spike. Bursting describes a tendency to fire short sequences of action potentials with ISIs less than ~ 5 ms. We first discuss refractoriness, then bursting.

Relative refractoriness

A popular model of neuronal firing that efficiently captures refractoriness is the integrate-and-fire neuron. In this model, the cell integrates synaptic currents in its dendrites and fires when the membrane potential exceeds a random threshold, after which the membrane potential resets. The integrate-and-fire neuron with constant current inputs is a renewal process (as there is no memory of the process beyond the last spike) and the ISI distribution is determined by the statistics of the random threshold (Gabbiani and Koch, 1998).

Figure 4.2 presents analysis of simulated data from an integrate-and-fire model with a random threshold. The ISI histogram (and therefore the threshold) is distributed according to a gamma function, order five. The spectrum is suppressed at low frequencies and the interval spectrum is constant. Spectral suppression at low frequency is often interpreted

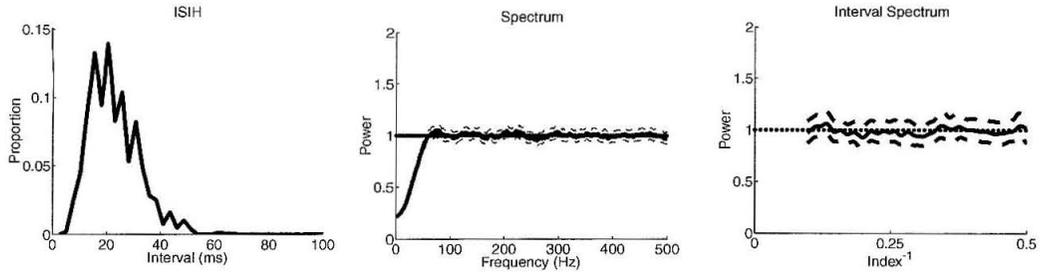


Figure 4.2: Activity from simulated renewal process. a) The ISI histogram. b) Spectrum normalized by the rate (solid). Poisson limit (dotted). c) The interval spectrum (solid) with 95% confidence intervals (dashed). The dotted horizontal line shows the spectrum of a renewal process for comparison.

as being a result of refractoriness (Gabbiani and Koch, 1998; Bair et al., 1994). Since refractoriness is defined by a time to the next spike, features in the spectrum can only be interpreted in this way if the activity comes from a renewal process. Since the interval spectrum is constant for this example, spectral suppression is due to refractoriness.

Bursting

Bursting can be simply included in the integrate-and-fire model by having the cell fire more than one spike each time the threshold is crossed. This is also known as the alternating renewal process. Figure 4.3 presents analysis of simulated data for a model where the cell fires two spikes 2 ms apart. The ISI histogram is the same as the previous model with an additional peak at short times due to the bursts. The spectrum contains one peak at low frequency in the gamma band and another at the bursting frequency. The interval spectrum is constant apart from a delta-function peak at 0.5 which reflects the lack of correlation between the lengths of the intervals, except at every other lag.

ISI correlations

A simple model that extends the integrate-and-fire model to include more general ISI correlations than the alternating renewal process is provided by the Markov renewal process. The process is characterized by a sequence of states determined by a Markov chain with a

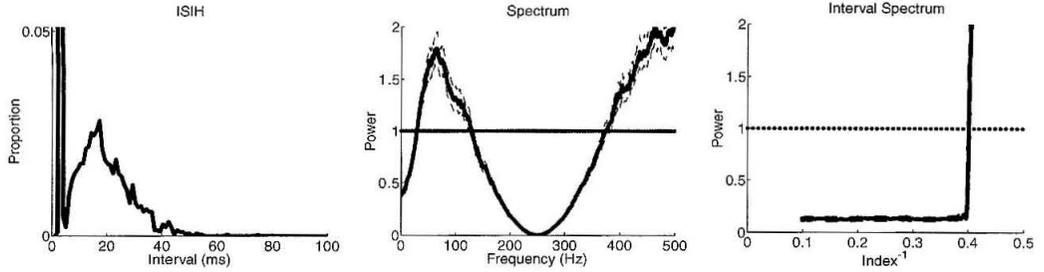


Figure 4.3: Activity from simulated process with bursting. a) The ISI histogram. b) Spectrum normalized by the rate (solid). Poisson limit (dotted). c) The interval spectrum for memory activity (solid) with 95% confidence intervals (dashed). The dotted horizontal line shows the spectrum of a renewal process for comparison.

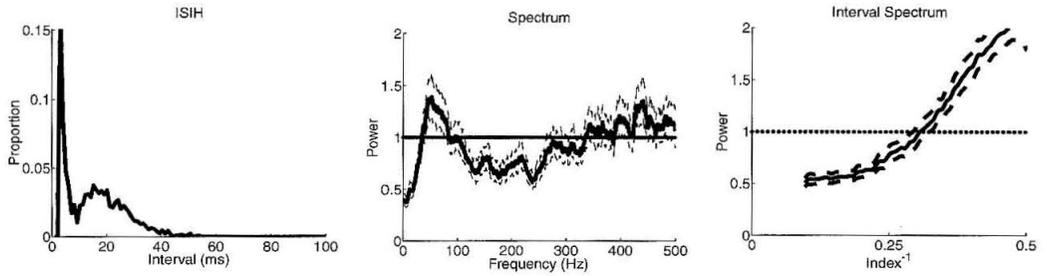


Figure 4.4: Activity from simulated Markov renewal process. a) The ISI histogram. b) Spectrum normalized by the rate (solid). Poisson limit (dotted). c) The interval spectrum for memory activity (solid) with 95% confidence intervals (dashed). The dotted horizontal line shows the spectrum of a renewal process for comparison.

matrix of transition probabilities, P , so that if the process is in state i , the probability is p_{ij} that the next state will be j (Cox and Isham, 1980). Since P is a matrix of probabilities, $\sum_j p_{ij} = 1$. Each state is a renewal process, intervals are independent and drawn from a distribution function, F_{ij} . The number of events in each state is geometrically distributed according to P .

We consider a Markov renewal process with two classes. The transition matrix is defined by $p_1 = 0.05$ and $p_2 = 0.5$. We assume the interval distribution depends only on the type of point at the beginning of the interval so $F_{11} \equiv F_{12} \equiv F_1$ and $F_{21} \equiv F_{22} \equiv F_2$. We model F_1 with an exponential function, rate 500 Hz, and F_2 as a gamma function, order 5, rate 50 Hz.

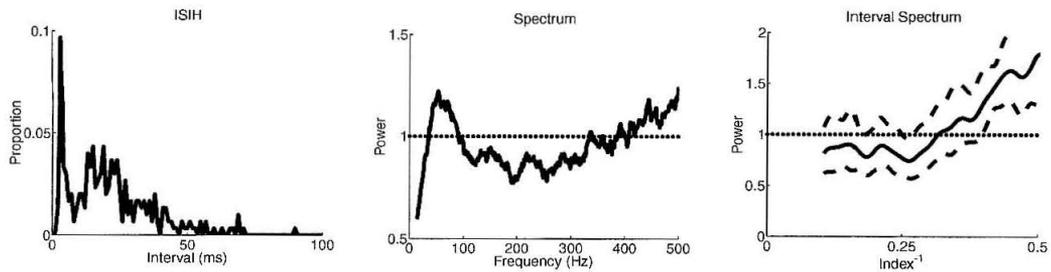


Figure 4.5: Activity from a sample LIP cell. a) The ISI histogram during working memory. b) The Spectrum normalized by the rate (solid). Poisson limit (dotted). c) The interval spectrum for memory activity (solid) with 95% confidence intervals (dashed). The dotted horizontal line shows the spectrum of a renewal process for comparison.

Figure 4.4 presents analysis of simulated data from this process. The bimodal ISI histogram results from the distributions from two classes. The spectrum is similar to Figure 4.3 and contains two peaks, one at low frequency, and the other at the bursting frequency. The interval spectrum, however, is quite different. This is related to the stochastic nature of the lifetimes of each state.

4.3.2 Interval analysis of cortical activity

ISI histograms and their relation to gamma band activity have been previously reported in macaque area MT (Bair et al., 1994) and in the visual cortex of the cat (Gray and McCormick, 1996). To determine whether there is additional temporal structure in cortical activity due to correlations between ISIs, we analyzed activity from three cortical fields: LIP, PRR and MT. We estimated the ISI histogram, the spectrum and the interval spectrum. The interval spectrum is analogous to the serial-order correlation function, but is a more sensitive measure that has not been previously presented for neuronal spiking activity.

Area LIP

We analyzed the intervals for a sample LIP cell. Figure 4.5a shows ISI histogram is bimodal: Intervals are either very short, < 3 ms, i.e., bursts, or center around 10–20 ms. Figure 4.5b shows the spectrum contains a gamma band peak and another peak at the

bursting frequency. Figure 4.5c presents the interval spectrum for memory activity from the sample cell on a log scale (solid) accompanied by 95% confidence intervals (dashed). The dotted line presents a uniform spectrum for comparison. The interval spectrum does contain significant structure with increasing energy at shorter lags. This means neighboring intervals are anti-correlated: short intervals are likely to follow long ones, and long intervals are likely to follow short ones. These observations are evidence that spiking activity in area LIP is non-renewal with single spikes and bursts that are temporally patterned in a broad gamma frequency band. The ISI correlations suggest a Markov renewal process, or some variation thereof, can model this activity.

Area MT

Previous analysis of temporal structure in MT reports gamma band peaks in the spectrum of many cells during the presentation of visual motion to the preferred direction (Bair et al., 1994). In that study, an investigation of the ISI histogram showed cells with gamma band peaks also burst. Since no evidence was found that the activity was not a renewal process, the authors concluded the temporal structure in the spectrum could be explained by structure in the ISIs, specifically the bursting. However, the interval spectrum was not estimated to test the renewal process model. Another model proposed by Bair et al. (1994) is that the cell fires bursts distributed according to a renewal process (cf. Chapter 4.3.1). Our analysis of ISI correlations explicitly tests these models.

Figure 4.6 presents analysis of activity from three typical cells during the presentation of highly coherent (51.2%–99.9%) visual motion to their preferred direction.

The first row, labelled Poisson, presents a cell that does not burst. It has a spectrum that is suppressed at low frequency but is otherwise constant and an interval spectrum that is constant. These features are consistent with a renewal process, and the activity is well-modeled by an integrate-and-fire neuron.

The second row, labelled Bursting, presents a cell that bursts. The ISI histogram has a narrow peak at short intervals and the spectrum has a peak at high frequency and rises up at low frequency. These features are due to many spikes in bursts arriving with the

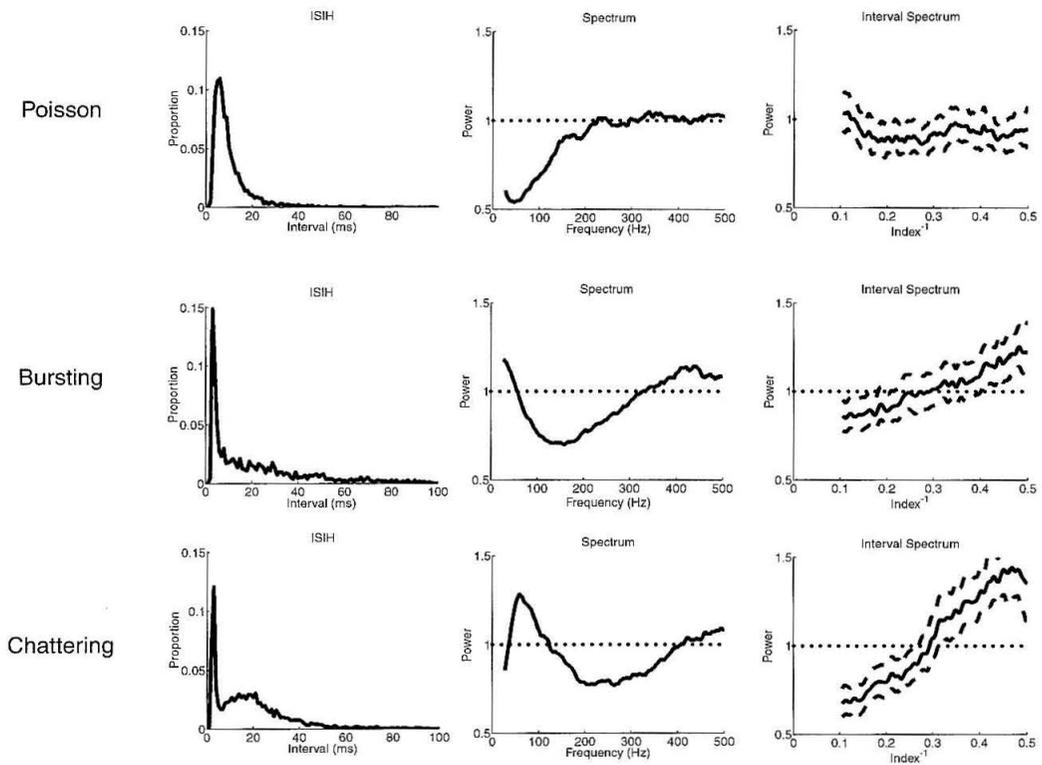


Figure 4.6: Activity from three sample MT cells. a) The ISI histogram. b) Spectrum normalized by the rate (solid). Poisson limit (dotted). c) The interval spectrum for memory activity (solid) with 95% confidence intervals (dashed). The dotted horizontal line shows the spectrum of a renewal process for comparison.

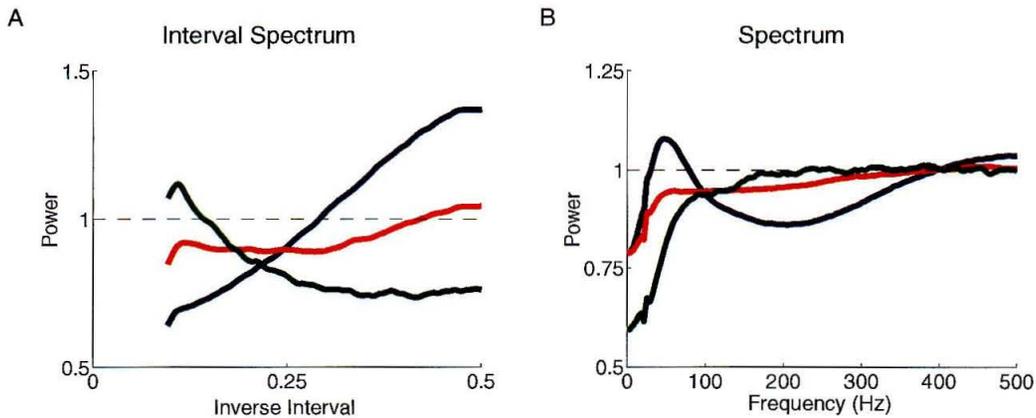


Figure 4.7: Activity from population of MT cells. a) The ISI histogram. b) Spectrum normalized by the rate (solid). Poisson limit (dotted). c) The interval spectrum for memory activity (solid) with 95% confidence intervals (dashed). The dotted horizontal line shows the spectrum of a renewal process for comparison.

same spacing. However, the interval spectrum is not constant. The positive slope means the lengths of intervals are anti-correlated and means the length of interburst intervals is correlated with the length of the bursts.

The third row, labelled Chattering, presents a cell that also bursts, as shown by the narrow peak in the ISIH. Like the Bursting cell, the spectrum has a peak at high frequency and rises up below 100 Hz. Unlike the bursting cell, the ISIH contains an additional broad peak at 20 ms and the spectrum drops to a value less than 1 at DC. The fall in the spectrum results in a prominent gamma band peak. The interval spectrum shows the intervals are anti-correlated. Bair et al. (1994) attribute the drop in the spectrum to refractoriness, but as the burst intervals are correlated with the length of longer "refractory" intervals, the dip in the spectrum at DC is not explained by a "time to last spike" phenomenon. An alternative explanation is that a driving process correlates the intraburst and interburst intervals leading to a Markov renewal process, or some variation thereof, similar to that found in LIP.

The analysis of sample cells showed that some cells exhibit activity that is not consistent with a renewal process. In particular, the departures best resolved with the interval spec-

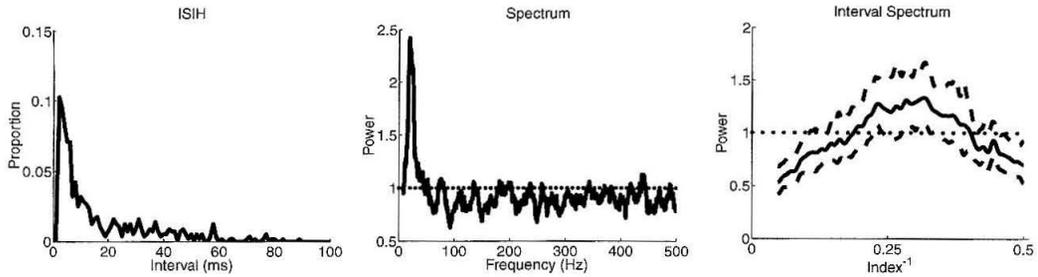


Figure 4.8: Activity from a sample PRR cell with gamma band spectral peak. a) The ISI histogram. b) Spectrum normalized by the rate (solid). Poisson limit (dotted). c) The interval spectrum for memory activity (solid) with 95% confidence intervals (dashed). The dotted horizontal line shows the spectrum of a renewal process for comparison.

trum were present over the whole range of the function. Therefore, we decided to divide the population of cells into three classes: those with an interval spectrum that was constant, had positive slope, or negative slope. Of the 174 cells in the study, 93 had no slope, 83 had positive slope and 8 had negative slope. We were interested to see if there were a difference in the average spectrum in each class. Figure 4.7a plots the average interval spectrum for each class of cells and Figure 4.7b, the average spectrum. The results show that cells without ISI correlations or some positive ISI correlations have spectra that are constant with some suppression at low frequency but no peak. In contrast, cells with negative ISI correlations have spectra with a lot of structure, particularly at high frequency and in the gamma band. These cells have ISI correlations that can be well captured by the Markov renewal process.

PRR

Recently, (Buneo et al.) report that the spectrum of spiking in PRR contains a gamma band spectral peak during a memory-reach task. Interestingly, across the population the spectral peak is at a slightly lower frequency than those reported for LIP and MT. Here the spectrum, ISI histogram and ISI spectrum are presented for these data.

Figure 4.8 presents the analysis for a sample PRR cell with gamma band activity. The ISI histogram shows many spikes arrive in bursts, similar to cells shown from LIP and MT,

but a number of differences exist. The ISI is not bimodal, the spectrum does not have a peak at the bursting frequency, and the interval spectrum does not have a peak at 0.5. The first two of these differences indicate bursting does not occur with as precise an ISI as in LIP and MT. The peak in the interval spectrum at $1/3$ indicates bursts arrive in triplets, not doublets. In combination with the observation that the firing rate is as high as is found in LIP (Snyder et al., 1997) but the spectral peak is at a lower frequency, indicates the cell fires more spikes per burst but with a longer interburst interval. Across the population of cells, 40 out of 87 showed these or similar features. Since the distribution of spike counts per state in the Markov renewal process is geometric, this structure is not well modelled by this process.

4.4 Discussion

Using the interval spectrum we have shown that ISIs are not independent in a substantial class of cortical cells. There is temporal structure due to correlations between intervals; therefore, the activity cannot be modeled by a renewal process. The Markov renewal process with two states models aspects of activity in LIP and MT, but does not capture activity from PRR.

Work suggests cells showing ISI correlations form a special class of cortical cells that are called "chattering" or "fast rhythmic bursting" (Gray and McCormick, 1996; Steriade et al., 1998). What are the characteristics of these cells? *In vitro* experiments on cortical slices shows that chattering in response to constant depolarizing currents is critically dependent on the interplay between afterhyperpolarization and afterdepolarization currents that follow each spike (Brumberg et al., 2000). These currents may mediate transitions between two states such as used in a Markov renewal process. However, this description seems to be incomplete, as it does not model ISI correlations observed in PRR. In this case, in place of a constant current input to the cell, we postulate the existence of another process that modulates that dependencies between intervals which results in more triplets than doublets.

Appendix A: Visual stimuli induce waves of electrical activity in turtle cortex

The computations involved in the processing of a visual scene invariably involve the interactions among neurons throughout all of visual cortex. One hypothesis is that the timing of neuronal activity, as well as the amplitude of activity, provides a means to encode features of objects. The experimental data from studies on cat support a view in which only synchronous (no phase lags) activity carries information about the visual scene (Gray et al., 1989). In contrast, theoretical studies suggest, on the one hand, the utility of multiple phases within a population of neurons as a means to encode independent visual features and, on the other hand, the likely existence of timing differences solely on the basis of network dynamics. Here we use wide-field imaging in conjunction with voltage sensitive dyes to record electrical activity from the virtually intact, unanesthetized turtle brain. Our data consist of single-trial measurements. We analyze our data in the frequency domain to isolate coherent events that lie in different frequency bands. Low frequency oscillations (< 5 Hz) are seen in both ongoing activity and activity induced by visual stimuli. These oscillations propagate parallel to the afferent input. Higher frequency activity, with spectral peaks near 10 and 20 Hz, is seen solely in response to stimulation. This activity consists of plane waves and spiral-like waves, as well as more complex patterns. The plane waves have an average phase gradient of $\sim \pi/2$ radians/mm and propagate orthogonally to the low frequency waves. Our results show that large-scale systematic differences in neuronal timing are present during visual processing.

Introduction

The flow of visual information through the retino-geniculo-cortical pathway in turtles projects to a rostral area of its dorsal cortex (Hall and Ebner, 1970; Desan, 1988), whose cellular and synaptic physiology share basic features with mammalian neocortex (Connors and Kriegstein, 1986; Kriegstein and Connors, 1986). Anatomical studies indicate that the geniculo-cortical afferents traverse the lateromedial width of dorsal cortex with few collaterals and their ending define the border between a lateral, solely visual area (D2) and a medial area (D1) that receives input from multiple sensory modalities (Desan, 1988; Heller and Ulinski, 1987) (Fig. A.1a). These afferents, along with widely distributed intrinsic fibers (Cosans and Ulinski, 1990), form an extensively interconnected lattice that is reminiscent of an association area. This organization is consistent with the finding that single neurons respond to stimuli in all quadrants of the visual field (Mazurskaya, 1972) and that visual stimulation leads to oscillations in the local field potential (LFP) throughout all of visual cortex (Prechtl, 1994). The frequency content of the oscillations is state dependent; a band near 20 Hz is specific to visual stimuli while lower frequency bands depend more generally on arousal (Prechtl and Bullock, 1995).

A.1 Methods

A.1.1 Preparation

Pond turtles (*Pseudemys scripta*) of both sexes with carapace lengths between 12 and 16 cm were maintained in aquaria (20° to 25° C). Use of the animals were under the guidelines of the Marine Biological Laboratory at Woods Hole and the National Institutes of Health. Under anesthesia by cold narcosis, the dorsal aspect of the telencephalon was exposed as described (Prechtl, 1994). The brain was partially isolated by sectioning the spinal cord and cranial nerves IV-XII from a ventral approach. The innominate artery was catheterized and the animal was continuously perfused with an oxygenated artificial cerebrospinal fluid (CSF); 96.5 mM NaCl, 2.6 mM KCl, 4.0 mM CaCl₂, 2.0 mM MgCl₂, 31.5 mM NaHCO₃, 10

mM dextrose and 1.5% (w/v) 60 - 90 kD dextran; saturated with 5% (v/v) CO_2 in O_2 , pH 7.4. The pial surface of the brain was stained for 15 min. with 0.02 - 0.08% (w/v) solutions of voltage-sensitive dye (JPW1114 or RH795; Molecular Probes, Eugene, OR) in CSF. The dye was observed to stain tissue throughout the entire depth of dorsal cortex (~ 500 μ m).

At the end of each experiment the animal was injected with a lethal dose of sodium pentobarbital; we recorded the response from the isoelectric brain as a measure of our noise level. Further, the exposed surface of the turtle cortex was labelled with pin holes at the center of the optical field; these marks and the LFP electrode locus served as fiducials for the location of the D1/D2 border. The brain was subsequently sectioned at 30 μ m and stained with cresyl violet; the D1/D2 border is defined by differences in cell packing and the width of the subcellular layer (Desan, 1988).

A.1.2 Optical recording

A 3.5 mm diameter region that encompassed most of rostral dorsal cortex was imaged with a 0.5 numerical aperture epi-illumination system (Precht and Bullock, 1995) onto an octagonal array of photodiodes whose individual outputs were band-limited (0.3 Hz single-pole high-pass and 300 Hz 4-pole Bessel low-pass filter) (Kleinfeld and Delaney, 1996) and digitized at $2f_{Nyquist} = 10$ Hz for 8 s epochs under the control of NeuroPlex acquisition software (Universal Imaging, West Chester, PA). The intensity of emitted light at time t for a pixel centered at (x, y) is denoted $I(x, y, t)$. The spatially averaged change in transmembrane potential within the field of each pixel is linearly proportional to the fractional change in emitted light, i.e., $DV(x, y, t) \propto DI(x, y, t)/I(x, y, t)$ where $DI(x, y, t) = I(x, y, t) - I(x, y)$ is the filtered signal and $I(x, y)$ is the DC value prior to acquisition of the sequence. Defective optical channels are ignored and the final data contain 447 spatial measurements.

A.1.3 Electrophysiology

Data from an epipial field electrode were acquired (1 Hz single-pole high-pass filter) concurrently with the optical images. The location of the electrode was marked by electro-

depositing iron from the electrode at the end of each experiment.

A.1.4 Stimuli

The loom was a white brick, 9 x 12 cm, that moved toward the contralateral eye on a 45° azimuth from midline. The total excursion was from 30 cm (10° subtended angle) to 5 cm (60°) at a speed of 8 cm/s. The step illumination was a red (620 nm) light emitting diode; the light was diffused by covering the eye with moistened tissue paper.

A.1.5 Data reduction

Five numerical procedures were performed with the data.

Spectral estimation:

For a given time series, denoted $V(t)$, the power spectrum, $S(f)$, was obtained with a direct multitaper estimate (Thomson, 1982) given by

$$S(f) = (1/K) \sum_{k=1}^K S_k(f) \quad (\text{A.1})$$

where

$$S_k(f) = \left| \sum_{t=1}^T e^{-i2\pi ft} w_k(t) V(t) \right|^2 \quad (\text{A.2})$$

is the discrete Fourier transform of $V(t)$ multiplied by the k^{th} Slepian taper, denoted $w_k(t)$. The tapers form an orthogonal basis and are characterized by a bandwidth parameter, denoted W ; there are $K + 1 = 2WT$ such functions that are concentrated spectrally in the full bandwidth $2W$, where T is the record length. The concentration reduces bias in the spectral estimate, while the average over spectral estimates reduces variance. The bandwidth is chosen to achieve smoothing while preserving relevant spectral structure.

Denoising:

The presence of shot-noise in the optical signal makes direct visualization of the space-time data difficult. As the desired signal is correlated in space and time, filtering can be used to improve the signal-to-noise ratio. Temporal filtering is appropriate since the frequency spectrum contains peaks; however, spatial frequency filtering will blur sharp spatial features that are present in the data. Superior results to conventional smoothing or filtering methods were obtained by truncating a space-time singular value decomposition (SVD) of the data $V(x, y, t)$. The standard SVD is applied to the two-dimensional data matrix obtained by collapsing the spatial dimensions into a single index, denoted \mathbf{x} , i.e., $V(x, y, t) \rightarrow V(\mathbf{x}, t)$. The decomposition is given by $V(\mathbf{x}, t) = \sum_n \lambda_n F_n(\mathbf{x}) G_n(t)$, where the spatial modes $F_n(\mathbf{x})$ and the temporal modes $G_n(t)$ each form an orthogonal basis and R is the rank of the matrix \mathbf{V} . The singular values λ_n are presumed to be in descending order. The data from an active brain was empirically found to be characterized by a few large singular values and a long tail of values; that from the pentobarbital perfused animal contained only the tail. This result can be quantitatively explained by assuming that the data for an active brain is a low rank matrix corrupted by shot noise (Sengupta and Mitra, 1999). We thus truncate the SVD to keep a number of modes R_{max} consistent with the estimated shot noise to obtain a denoised matrix $V_{Denoised}(\mathbf{x}, t) = \sum_{n=1}^{R_{max}} \lambda_n F_n(\mathbf{x}) G_n(t)$. In practice, $R = 447$ and $R_{max} < 40$.

Temporal filtering:

Projection filters based on Slepian functions provide optimal frequency isolation and were used for purposes of temporal filtering. For a given length of time series, i.e., $V(t)$ with $t \in [1, T]$, the projection operator \mathbf{P} has matrix entries $P_{k,t} = e^{-2\pi i f t} w_k(t)$. The Slepian functions $w_k(t)$ are characterized by a bandwidth W (see above); the choice $k < 2WT$ projects the spectral energy into a frequency band $[f - W, f + W]$. The projected or filtered time series is given by $\mathbf{V}_{filtered} = \mathbf{P}^T \mathbf{P} \mathbf{V}$.

Spatial coherence:

To examine the spatial patterns of coherence corresponding to peaks in the frequency spectra, a space-frequency SVD (Mann and Park, 1994) was used. In this method, the space-time data is first projected into a local temporal frequency domain using the projection operator defined above.

$$\tilde{\mathbf{V}}_k(\mathbf{x}, t) = \sum_{t=1}^T V(\mathbf{x}, t) e^{2\pi i f t} w_k(t) \quad (\text{A.3})$$

Here the index k effectively defines a local frequency index in the band $[f - W, f + W]$ with $k < 2WT$. For a fixed frequency, f , a SVD is performed on the complex matrix $\tilde{\mathbf{V}}(\mathbf{x}, t)$ to yield

$$\tilde{\mathbf{V}}_k(\mathbf{x}, t) = \sum_{n=1}^R \lambda_n(f) F_n(\mathbf{x}; f) G_n(k; f) \quad (\text{A.4})$$

where R is the rank of $\tilde{\mathbf{V}}_k(\mathbf{x}, t)$, typically set by $2WT$. If the fluctuations were completely coherent in the given frequency band, only one of the singular values would be non-zero. A measure of coherence is then given by the ratio of the power of the leading mode to the total power.

$$C(f) = \lambda_1^2(f) / \sum_{k=1}^K \lambda_k^2(f) \quad (\text{A.5})$$

for a completely coherent response $C(f) = 1$ while for a uniform random process $C(f) = 1/K$. Where $C(f)$ has a peak, it is useful to examine the largest spatial mode $\tilde{\mathbf{V}}_1(\mathbf{x}, f)$. The magnitude of this complex image gives the spatial distribution of coherence at that frequency, while gradients in its phase indicate the local direction of propagation.

Demodulation:

Slow changes in the magnitude and phase of the oscillatory response near a particular frequency are found by complex demodulation of the denoised data. We form $\tilde{\mathbf{V}}(x, y, t) =$

$e^{2\pi ift}V(x, y, t)$. , so that in temporal frequency space the origin is now shifted to the center frequency f and low-pass filter each time series.

A.2 Results

A.2.1 Basic response

We measured the change in fluorescence from cortex stained with voltage sensitive dye (Precht and Bullock, 1995; Orbach and Cohen, 1983) together with measurements of the LFP (Fig. A.1a). The optical measurements provide a means to record the average transmembrane potential at ~ 450 neighboring sites with an intrinsic spatial resolution of ~ 150 μm per site (Orbach and Cohen, 1983). The optical signal at a single location atop dorsal cortex shows a large depolarization upon onset of a loom (Fig. A.1b), a behaviorally relevant stimulus (Killackey et al., 1972), consistent with previous results (Senseman, 1996). Superimposed on the depolarization are relatively fast oscillations (Fig. A.1b) whose timing is correlated with those seen in the LFP sampled at an adjacent location (Fig. A.1c; correlation coefficient = 0.6). We detect no changes in the intrinsic optical properties of the preparation upon stimulation (data not shown).

Both the LFP and the optical signal show weak spectral features prior to stimulation (Figs. A.1d,e). In the presence of a visual stimulus, the power is enhanced at all frequencies and a significant peak emerges near 18 Hz (* in Figs. A.1d,e). The average spectrum across all strongly activated sites in the optical field (Figs. A.2a,b) is similar to that seen at the previous single location (Fig. A.2e); a band centered at 18 Hz appears within 0.5 s of the onset of stimulation and lasts for nearly 2 s of the 3 s stimulation period (Fig. A.2a). Similar results are observed in all preparations with the loom ($n = 7$).

We tested if our preparation procedures distorts the spectral content of the signals. First, we compared the stimulus-induced LFP from both intact preparations and preparations in which the brain was isolated from nociceptive and discriminative sensory inputs. The spectral content of the LFP was essentially unchanged. Further, we observed that the spectral content was unchanged by vascular perfusion, a procedure that removed a system-

Figure 1

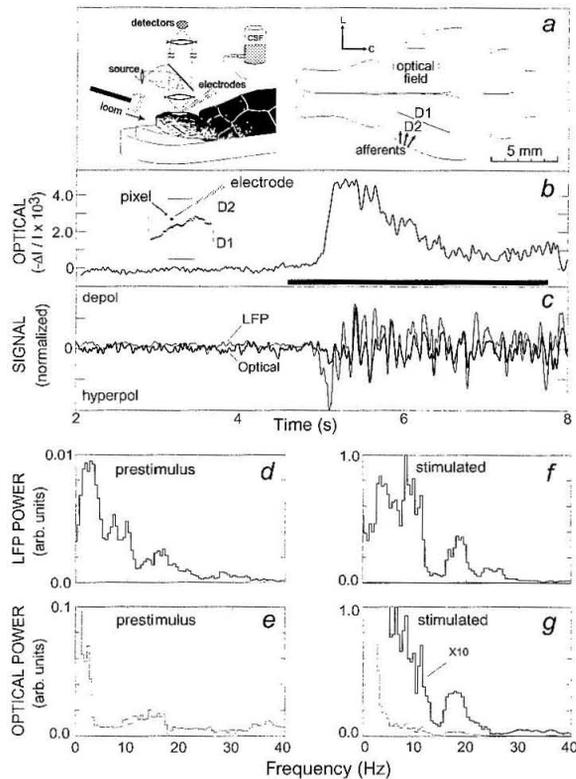


Figure A.1: Experimental procedure (a) and the response from single cortical locations (b-e). (a) Schematic of the set-up. A 3.5 mm diameter region of stained cortex that encompasses parts of lateral (D2) and medial (D1) dorsal cortex was imaged onto an array of photodiodes (135 μm /pixel). A metal electrode was used to measure the LFP. (b) The optical signal, low pass filtered at 40 Hz. The large depolarization appears ~ 400 ms after stimulus onset of the loom ($t = 0$); its decay rate is twice that for the high-pass filter. The insert shows the approximate location of the pixel relative to the reconstructed D1/D2 border (Sec. A.1). (c) Overlay of the optical signal at one pixel and the adjacent local field potential (LFP), as indicated in the insert in part b. The large depolarization in the optical signal was removed with a median filter (400 ms wide) and the relative amplitudes of the two signals were adjusted to give the maximal overlap during the stimulus interval. (d) Spectral power (logarithmic scale) in the LFP during the 1 to 2 s epoch prior to stimulation and the 1 to 2 s epoch after the onset of stimulation (2WT = 5.0; K=3). Note the emergence of a broad peak at 18 Hz (*) during stimulation. (e) Spectral power in the optical signal during the 1 to 2 s epoch prior to stimulation and the 1 to 2 s epoch after the onset.

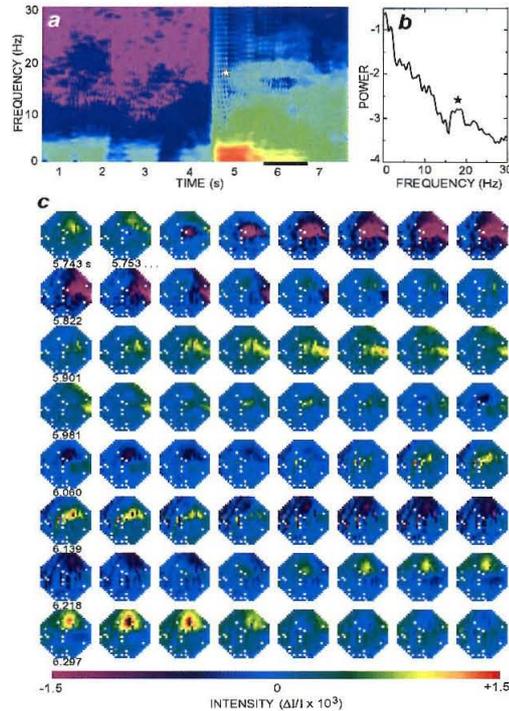


Figure A.2: The spatiotemporal optical signal. (a) The average of the power spectra for 72 pixels over the most active part of cortex, computed over sliding 1 s intervals ($2WT = 4.0$; $K = 2$). The logarithm of the data is false colored, with red corresponding to the maximum and purple corresponding to minimum values. The transition near 0.3 s reflects the onset of stimulus-induced activity. The star indicates the stimulus-induced band at 18 Hz. (b) The spectrum for a particular 1 s interval (bar in part a). (c) The broadband filtered optical data. We plot every 8th frame (126 Hz) of the single-trial response. The data were spatially denoised, low-pass filtered at 60 Hz, and the large stimulus-induced depolarization was removed with a median filter (400 ms width). The color scale for each frame was separately normalized. Note the net flow of depolarization in the rostral-to-caudal and medial-to-lateral directions. Orientation and size of the optical field as in Figure A.1a.

atic source of optical noise. Lastly, we checked for pharmacological effects of the voltage sensitive dye. Dye concentrations of 0.08% (w/v) or more decreased the magnitude of the "20 Hz" band in the LFP by 10%; concentrations at twice this level had a more pronounced effect. Thus our images may contain suppressed high frequency activity, although other sources of distortion, such as anesthetics or anticholinergic paralytics, are absent from the preparation.

A.2.2 Large-scale dynamics

Are the stimulus-induced oscillations in different regions of turtle visual cortex coherent with each other? The "raw" image data was denoised (Sec. A.1) and a front of depolarization (Fig. A.1b) that propagates along the direction of the afferents (Senseman, 1996) was subtracted. The resultant sequence of images exhibits spatial patterns that evolve over time. For the sequence of Figure A.2c, we observe depolarization and hyperpolarization that travel in the caudal direction, consistent with low resolution LFP measurements (Precht and Bullock, 1995), as well as swirls of polarization. This examination shows that the spatial and temporal aspects of the response are coupled. However, the images appear complex, not surprisingly in light of the broad spectral content of the response, and are dominated by the larger amplitudes of the low frequency activity.

As a means to probe for a simple spatiotemporal structure that might underlie the response, we determined if the observed cortical signal exhibits a dominant pattern of spatial coherence within separate frequency bands. We considered the response over 3 s intervals prior to and during stimulation (Sec. A.1); the magnitude of calculated coherence for each pixel in the dominant pattern reports the stability of the oscillation and the phase reports the relative timing of the oscillation during the 3 s interval. Prior to stimulation, significant coherence is limited to frequencies less than ~ 5 Hz (Fig. A.3a). The magnitude of the coherence is dispersed (false color plot, Fig. A.3b), but the phase gradient lies in the laterorostral to mediocaudal direction (contour plot, Fig. A.3b); this axis parallels the thalamocortical input tract (Figs. A.1a,b) (Ulinski, 1986). The presence of a stimulus causes the spatial coherence at low frequencies to apparently coalesce (Fig. A.3c) while the phase

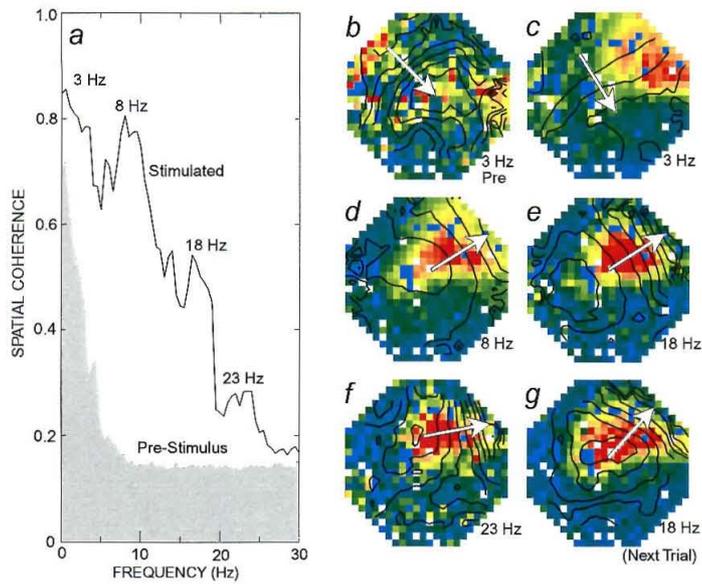


Figure A.3: The magnitude and phase of the spatial coherence of the optical signal from a space-frequency SVD analysis. (a) The coherence averaged over a $T = 3$ s interval both prior to and subsequent to the onset of stimulation. In the latter case the interval encompasses the entire epoch of significant spectral power in the band at 18 Hz (Fig. A.2a). The coherence was estimated at successive frequency bins ($2WT = 3.0$; $K = 7$); a value of $C(f)$ of 0.14 indicates the lack of significant coherence (b) Phase (contour lines) and amplitude (false color) of the coherence at $f = 3$ Hz prior to stimulation. The magnitude defines the spatial distribution of coherence while the phase defines the temporal delays between different regions. The relative magnitude is false colored with red for maximum and blue/green for zero and the phase is overlaid as a contour plot with $\pi/12$ radians per contour. The arrow indicates the dominant direction of the gradient. (c - f) Phase and amplitude at $f = 3, 8, 18,$ and 22 Hz, respectively, during stimulation by a loom. (g) Phase and amplitude at 18 Hz for the next trial with the same animal. Orientation of the optical field as in Figure A.1a).

gradient remains congruent to that seen prior to stimulation (cf. Figs. A.3b,c). Further, the stimulus elicits spatial coherence in bands centered near the higher frequencies of 8, 18, and 23 Hz (Fig. A.3a). The magnitude of the coherence for each of these additional bands are colocalized. Their phase gradients lie in the medio-rostral to latero-caudal direction (Figs. A.3d,e,f); this axis parallels the D1/D2 border and is perpendicular to the gradient of the ongoing low frequency activity (cf. Figs. A.2c,e). Repetition of the stimulus leads to a dominant pattern with a similar spatial localization and phase gradient (cf. Figs. A.2e,g).

The three essential features of the response discussed above were seen in all trials (7 animals, 3–4 trials/animal), vis, (i) the low frequency activity propagates predominantly along the afferent fiber tract, (ii) the high frequency activity propagates predominantly orthogonal to that at low frequency, and (iii) stimulus-induced coherent activity is localized. The variability in the direction of propagation for the dominant component was 25° (SD) for frequencies below ~ 5 Hz and 15° for frequencies ~ 8 Hz; for the 18 Hz band the variability drops to 5° . The direction of this phase gradient did not significantly change when the axis of the loom was shifted by 90° in the horizontal plane or when the oscillations were induced by a step of illumination rather than a loom ($n = 3$); thus the direction appears to be intrinsic to cortex. The gradient across the region of highest spatial coherence is 1.4 ± 0.2 radians/mm (mean \pm SD) for frequencies between 8 and 12 Hz and 1.6 ± 0.3 radians/mm for frequencies between 16 and 20 Hz. The inferred propagation speeds are 0.05 and 0.09 mm/ms, respectively; such slow speeds are reported for fibers in turtle cortex.

Spatial events that are short-lived or whose pattern drifts will not survive the averaging used to determine the dominant pattern (Fig. A.3), yet may be fundamental to the underlying dynamics. We focus on the temporal evolution of the stimulus induced response within the 18 Hz band; demodulation (Sec. A.1) of the data provides a means to isolate the magnitude (color saturation; Fig. A.4) and phase (color hue; Fig. A.4) of the response. The phase is poorly defined at the onset of stimulation and shows an initial coherent response that corresponds to linear waves (* in Fig. A.4a) that run in the caudolateral direction, as was seen in the dominant mode of the coherence (Fig. A.3e). At later times we observe the onset of a phase singularity and circular waves (** in Fig. A.4a), followed by a return to linear waves (***) in Fig. A.4a). The singularity lies near the D1/D2 border but does not correspond to any presently known anatomical feature or to injury. The epicenter wanders as the response progresses; this partly accounts for the absence of a circular component in the dominant mode over the 3 s stimulation period, described above (Fig. A.3e). Lastly, the presence of a circular wave at 18 Hz persists when the demodulate is averaged over independent time intervals (4 cycles; Fig. A.4b). Consistent with a lifetime of at least 4 periods, the circular wave at 18 Hz is not accompanied by similar waves at neighboring

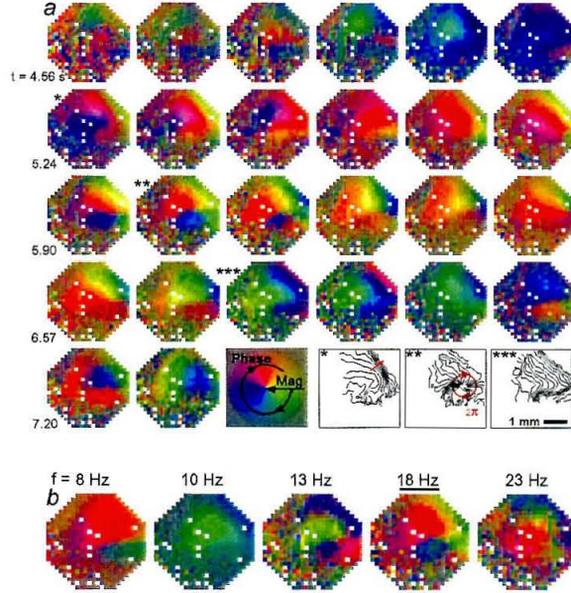


Figure A.4: Temporal evolution of the band-limited electrical activity. (a) The magnitude and phase of the optical response centered at 18 Hz as a function of time, beginning with the onset of the loom. We plot the demodulate, $\tilde{V}_f(x, y, t)$, so that color-saturation level is normalized by the magnitude of the demodulate and the color hue codes the phase. Note the presence of approximately linear (*) and circular (**) phase shifts. Each frame corresponds to an independent sample (9 Hz); a phase shift of 2π radians corresponds to a cycling through the chromatic scale (e.g., red to green to blue to red). The final row includes contour maps of the phase only at times $t = 0.68$ s (*), $t = 1.45$ s (**), and $t = 2.22$ s (***); each contour corresponds to $\pi/12$ radians. Orientation of the optical field as in Figure A.1a). (b) The magnitude and phase of the demodulate centered at $f = 8, 10, 13, 18,$ and 23 Hz, filtered at $W = f/4$ ($2WT=2$; $K=1$) and averaged over a period of 4 cycles. The center time of the interval was $t = 1.51$ s (** in part a).

frequencies (14 and 23 Hz; Fig. A.4b).

A.3 Discussion

The observed phase gradient for the high frequency bands is likely to be mediated by intrinsic cortical connections since it lies orthogonal to the geniculo-cortical afferents. However, the known short- and long-range horizontal connections exhibit no obvious orientation (Cosans and Ulinski, 1990). Thus the direction of the gradient at high frequencies is likely to emerge from the dynamics of the underlying neuronal oscillators and their connections (Grannan et al., 1993; Ermentrout and Kopell, 1994).

The phase gradients seen here (Figs. A.3, A.4) and those reported for other central structures (23-32) do not exceed 2π radians. This suggests that the oscillatory component of the spike train of different neurons in cortex have an unambiguous phase. Gradients with unambiguous phase provide a unique label that may be used to segment (von der Malsberg and Schneider, 1986) or categorize (Sompolinsky and Tsodyks, 1994; Hopfield, 1995) the visual scene in the temporal domain.

The wave phenomena reported here are likely to modulate the computations performed by the turtle's visual system, although there is presently no evidence to relate such spatiotemporal patterns to behavior in a causal fashion. A similar situation holds for the visually-induced slow (~ 0.1 Hz) waves of heat that propagate across rat cortex (Bouyer et al., 1981) these report changes in metabolism and blood flow that are secondary to increases in neuronal activity, and for the visually-induced "40 Hz" oscillations in the electrical activity of cat and monkey cortex (Eckhorn et al., 1988; Gray et al., 1989; Kreiter and Singer, 1996). While the latter oscillations are believed to be spatially uniform, high resolution imaging may reveal features similar to those reported here (Figs. A.3, A.4).

Appendix B: The role of nonlinear dynamics of the syrinx in the vocalizations of a songbird

B.1 Introduction

Song is characterized by the modulation of sound properties over a wide range of timescales (Greenewalt, 1968). Understanding the mechanisms by which the brain organizes this complex temporal behavior is a central motivation in the study of the song control and learning system (Nottebohm et al., 1976, 1982; McCasland, 1987; Scharff and Nottebohm, 1991; Konishi, 1994; Vu et al., 1994; Yu and D., 1996). Here we present evidence that, in addition to central neural control, a further level of temporal organization is provided by nonlinear oscillatory dynamics intrinsic to the avian vocal organ. A detailed temporal and spectral examination of zebra finch (*Taeniopygia guttata*) song reveals a class of rapid song modulations consistent with transitions in the dynamical state of the syrinx. Furthermore, in vitro experiments demonstrate that the syrinx can produce a spectrally and temporally complex sequence of oscillatory states in response to the slow variation of respiratory or syringeal parameters. As a consequence, simple variations in a small number of neural signals can result in a complex acoustic sequence.

B.2 Methods

B.2.1 Isolated syrinx

Adult male zebra finches (Canary Bird Farm, Old Bridge, NJ) were anesthetized and sacrificed by decapitation. The syrinx was removed, cleaned of excess connective tissue, and mounted in the experimental chamber. The trachea was connected to an adjustable vacuum source, and one bronchial tube was connected to a source of humid air with adjustable

pressure. The other bronchial tube was pinched closed and sealed with cyanoacrylate. The trachea and anterior portion of the syrxinx was embedded in 3% agarose in 0.1M phosphate buffer. Tracheal and bronchial pressures were recorded with electronic pressure sensors (Sensym, Inc.). Above a threshold flow rate (~ 0.5 standard liters per minute), the labium internum (LI) and the medial tympaniform membrane (MTM) spontaneously oscillated over a wide range of flow rates (0.5–2.5 SLPM) and pressure differentials between the trachea and bronchus (0.1–1.0 kPa). Oscillation at pressures or flow rates higher than this physiological range resulted in rupture of the medial membranes. Acoustic signals generated by the syrxinx were recorded using either an electret microphone placed roughly 2 cm from the syrxinx, or with a probe microphone placed into the trachea.

B.2.2 Syrxinx response to parameter changes

Dependence of the oscillatory state of the excised syrxinx on parameters was examined for variations in both tracheal pressure and membrane tension. To study pressure variations, the bronchial inlet was held at a constant pressure (0–0.5 kPa) while the tracheal outlet pressure was smoothly ramped down 1–2 kPa below atmospheric pressure, and then ramped back to atmospheric pressure. The control pressures were measured ~ 50 cm from the syrxinx to reduce the sensitivity of the pressure measurement to changes in syringeal resistance. The dependence of oscillatory state on the position of the external labium (LE) was examined during sustained syringeal oscillation at a constant tracheal pressure. The LE was manually forced with forceps into the bronchial lumen until the oscillation ceased. Simultaneous video recording of the preparation and the acoustic signal permitted later quantification of LE displacement and the associated acoustic modulations. Displacement of the LE produced no apparent displacement or tensioning of the LI or MTM, so the frequency changes were probably a result purely of constriction of the bronchial lumen, rather than tension changes.

B.2.3 Stroboscopic imaging

Stroboscopic images of the oscillating syrinx were obtained in vitro to examine in detail the motion of the labium internum (LI) and the medial tympaniform membrane (MTM). The syrinx was induced to oscillate as described above. Pulsed illumination (10 ms duration) from a cluster of four bright LEDs was synchronized with the fundamental oscillation period of the membrane as follows. The measured acoustic signal was rectified, band-pass filtered around the oscillation frequency (typically 500–1000 Hz), and used to trigger a variable-delay pulse generator. The pulse generator triggered the LEDs at a known time relative to the acoustic signal. The delay was slowly swept through several milliseconds to sample all phases of the oscillation; the images were recorded with a CCD camera and video recorded with a Sony Betacam recorder for later analysis.

B.2.4 Spectral analysis

Spectrograms were computed with direct multitaper spectral estimates²¹ with time bandwidth product $TW=2$. A sliding window of length 1 ms was employed. In Figure B.1c, superimposed on the spectrogram is the F-spectrum (Thomson, 1982). This quantity measures the goodness-of-fit to a sinusoid at any given frequency. The F-spectrum was computed with a moving window as above, and the results were thresholded at a fixed confidence level before being added to the displayed spectrogram.

B.3 Results

B.3.1 Analysis of bird song

Zebra finch songs are organized into short segments of sound, referred to as syllables, that often contain rapid sequences of whistled notes, harmonic stacks, and aperiodic signals (Immelman, 1969; Price, 1979). Sounds within a syllable are produced contiguously with no silent interval during the transitions between distinct signal types (Figure B.1a). This continuity allows one to observe changes in the oscillatory state of the syrinx during modu-

lations of the acoustic signal. We recorded the directed song of male zebra finches ($n=12$) in a sound isolation chamber. Close examination of the song spectra and acoustic waveforms shows sudden transitions from periodic to aperiodic or chaotic dynamics, period doubling, and mode-locking transitions. We believe these features arise because the syrinx behaves as a low-dimensional nonlinear dynamical system.

Nearly all (10 of 12) of the songs we recorded contained syllables exhibiting rapid (< 10 ms) transitions from periodic to chaotic, noisy, signals. Figure B.1(b) shows two examples of this behavior at an expanded time scale. These transitions sometimes occur in less than 1 ms, and often appear to jump back and forth between the periodic and chaotic oscillation. Note that the transition from a nearly periodic signal to a noisy signal occurs in roughly 1 ms. Another characteristic of nonlinear dynamical systems, period doubling, is characterized by a change in oscillation frequency such that spectral components appear at half the original frequency spacing. An example of period doubling in zebra finch song is shown in Figure B.1(c); at the time marked with the (*), note the rapid onset of interspersed harmonic components as described above. A rapid transition in oscillation frequency occurs at the point marked (2-3) such that the 3rd harmonic after the transition has approximately the same frequency as the 2nd harmonic prior to the transition. Roughly 15 ms later, another transition occurs such that the 4th harmonic is suddenly shifted down to the frequency of the third harmonic. Period doubling transitions were seen in three of the zebra finch songs examined.

Zebra finch song syllables can also exhibit modulations suggesting the presence of mode locking in the syringeal dynamics. Mode locking occurs when two components of a system are constrained by some nonlinear interaction to maintain a small integer ratio of oscillation frequencies. If the characteristic frequency of one component of the system is changed relative to the other, mode locking transitions may occur as the oscillation frequency suddenly jumps to achieve a new stable integer ratio with the fixed frequency (Fletcher, 1978; Jackson, 1992). In zebra finch song, for example, syllables often exhibit an overall downward sweep in fundamental frequency (Price, 1979). However, this downward frequency sweep may be punctuated by steps, or rapid transitions in fundamental frequency (Fig. B.1c). In

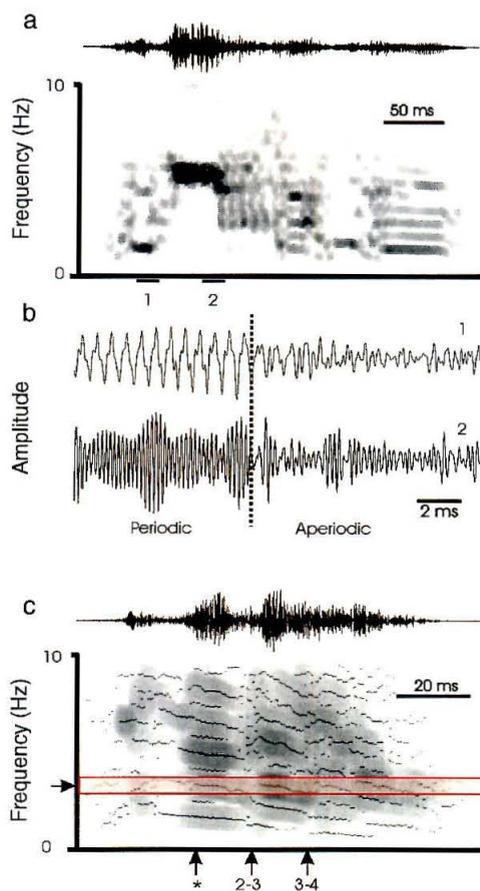


Figure B.1: Spectral and temporal characteristics of zebra finch song. a) The spectrogram of a syllable displaying the wide variety of spectral features common in zebra finch song. The pressure waveform is shown at the top. b) The rapid transitions between these signals are seen by expanding the pressure waveform around the regions marked 1 and 2 in (a). c) Spectrogram and harmonic analysis of a syllable recorded from a different bird. Note the rapid frequency doubling transition at the time marked with the (*). The 2nd, 3rd and 4th harmonic in this sequence are all roughly aligned with the bright spectral band at 3200 Hz, indicated by the horizontal bar and arrow.

the example shown, the mode locking transitions occur such that the oscillation frequency maintains a ratio of two, three and four with respect to a common frequency at roughly 3 kHz. Similar features were observed in 9 of the 12 zebra finch songs studied.

B.3.2 Mechanism of song generation

We wish to distinguish between two hypotheses for the origin of the song features described above: direct muscular manipulation during singing of the syringeal parameters by the song control system, or alternatively, intrinsic dynamics of the syrinx. To this end, we examined the acoustic signals generated by an in vitro preparation of the zebra finch syrinx ($n=6$). The excised syrinx was induced to oscillate by applying low air pressure at the trachea and high pressure at the bronchial inlet. In vitro oscillations of the zebra finch syrinx were quite similar in oscillation frequency and amplitude to normal zebra finch song and occurred over a physiologically reasonable range of pressures and flow rates (Hartley and Suthers, 1989). Spectra of the sound signals measured outside the syrinx (Fig. B.2a) show an enhancement of lower harmonics compared to normal song spectra. The high harmonic components in song are thought to be produced by the rapid pressure fluctuations associated with closure of the bronchial lumen during syringeal oscillation; we expect that such rapid fluctuations would not be present in the sound signal measured outside the syrinx.

Note that as a function of LE displacement, the fundamental frequency exhibits a series of discrete jumps such that the third harmonic, second harmonic, and fundamental are roughly aligned with a common frequency at ~ 2.2 kHz.

A smooth ramp-down and ramp-up of air pressure applied at the tracheal outlet produced a sequence of oscillatory states, typically with jumps in oscillation frequency and, at higher pressures, transitions from periodic to aperiodic oscillation (Fig. B.2a). The precise sequence of acoustic signals produced during this manipulation was different for each preparation. The acoustic sequence was also sensitive to small changes in various experimental parameters, such as tension on the bronchus. Intrinsic differences in syringes and sensitivity to parameters may partially account, respectively, for song variations across individuals and variability within individuals. Another feature commonly observed was hysteresis in the

acoustic signal; the sequence of acoustic signals generated as the pressure was ramped up usually differed in a repeatable manner from the sequence as the pressure was ramped down. The transitions observed during this manipulation were usually abrupt, often occurring in roughly one period of the acoustic signal (Fig. B.2b).

Further evidence for nonlinear dynamics of the syrinx was seen during simple manipulations of syringeal configuration. Deflection of the labium externum (LE) (see Fig. B.3) into the bronchial lumen constricts the flow of air through the syrinx, and is thought to be a normal articulation during song production. In all of the syringes tested ($n=3$), this constriction resulted in a large increase in the fundamental oscillation frequency during in-vitro oscillation. In two syringes, however, the frequency was not a smooth function of LE displacement; rather the frequency exhibited sudden jumps at certain discrete displacements (Fig. B.2c). Furthermore, the oscillation appears to be constrained to frequencies that are close to a small integer sub-harmonic of a common (higher) frequency. In the example shown, the frequency jumps from ~ 700 Hz to ~ 1100 Hz to ~ 2200 Hz, which have periods, approximately, of three, two, and one times the period of the common 2200 Hz component.

To understand the origin of the apparent strong nonlinearity in the syrinx, we used a stroboscopic imaging technique to visualize in vitro oscillations of the zebra finch syrinx ($n=5$) (Fig. B.3). In vitro, there was a threshold flow rate (~ 0.5 standard liters per minute) at which the labium internum (LI) and the medial tympaniform membrane (MTM) began to oscillate with large amplitude motion (~ 0.5 mm). The motion of these membranes was phase locked to each other and to the generated sound, suggesting that these membranes are in fact the source of the sound. In all syringes studied, the anatomical structure exhibiting the largest motion amplitude was the LI, rather than the much lighter MTM, in agreement with recent evidence that the LI plays a central role in syrinx function (Goller and Larsen, 1997). Close examination of the strobed oscillation under a microscope indicated that the bronchial lumen is probably briefly closed by contact of the internal labium with the external labium during part of the oscillatory cycle; however, the wave-like motion of the medial membranes made it difficult to determine quantitatively the fraction of time during the

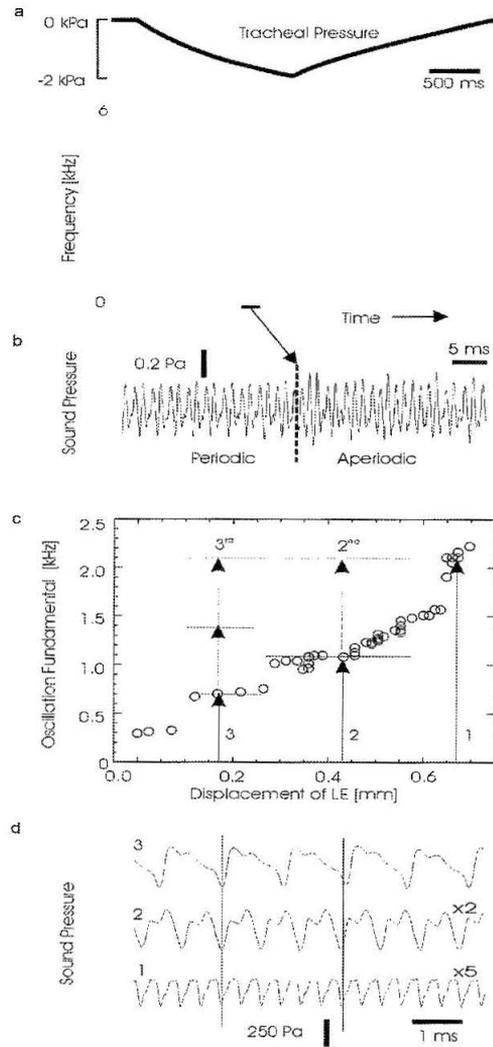


Figure B.2: Patterns of acoustic signals generated in the in-vitro syringe preparation. a) As the tracheal pressure (top) is varied over a range of 2 kPa, the sound generated by syringe exhibits a sequence of modulations. Bronchial pressure is held constant (0.2 kPa). b) The pressure signal shows the transition occurs in roughly one oscillatory period. c) Displacement of the external labium (LE), into the bronchial lumen, increases the fundamental oscillation frequency before the oscillation ceases. d) Acoustic waveform measured at labial displacements indicated in (c). The three waveform segments are shown at different scales, as indicated.

oscillation that the bronchial lumen was closed. Because the excursions of these membranes are comparable to the size of the bronchial lumen, the Bernoulli forces driving the oscillation (Fletcher, 1988) are a highly nonlinear function of membrane displacement.

B.3.3 Syringeal modeling

To determine if our stroboscopic and acoustic observations of *in vitro* syringeal oscillations were consistent with a view of the syrinx as a nonlinear dynamical system, we developed a biophysical model of the syrinx. We followed previous work on both human (Ishizaka and J.L., 1972) and avian (Fletcher, 1988) vocalizations, as well as models of fluid flow through collapsible tubes (Bertram and Pedley, 1982). The model is constituted by a set of ordinary differential equations describing the mechanical motion of the syringeal membranes in the presence of airflow through the bronchial tube. The essential nonlinearity in the model is the nonlinear pressure-flow relation associated with the Bernoulli force. The membranes were modeled as a two mass system. Membrane parameters were based on estimates of the resonant mode structure of the LI and the MTM. The essential result is that the model exhibits qualitatively similar behavior to *in-vivo* zebra finch song as well as the *in-vitro* syringeal oscillations. The model shows period doubling and transitions from periodic to chaotic dynamics (Fig. B.4), as well as mode locking transitions (not shown). Our model results suggest a mechanism for the observed mode-locking, namely coupling of the basic Bernoulli oscillation to a higher vibrational mode in the membranes. An additional mechanism underlying mode-locking might be coupling of the oscillatory mechanism to tracheal resonances.

B.4 Discussion

The generic signatures of nonlinear dynamics described here have also been described in models of human vocalizations (Berry et al., 1994), and in patients with vocal disorders (Hertzel et al., 1994). A distinction with human vocalizations demonstrated by our model and the *in vitro* results is that detailed spectral shaping in zebra finch song arises in large

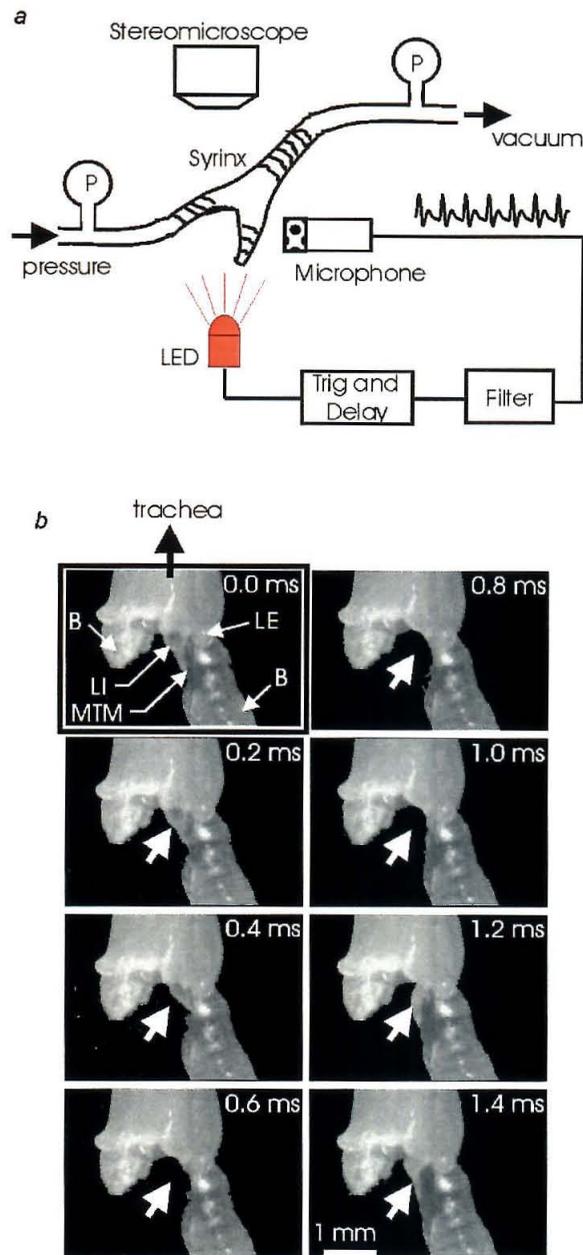


Figure B.3: Stroboscopic observation of in-vitro syringeal oscillations. a) Experimental setup. The syringe is induced to oscillate in vitro by applying suction to the trachea and slight overpressure to one bronchial tube. The sound pressure measured ~ 2 cm from the syringe is used to trigger an array of bright LEDs with a variable time delay. b) The configuration of the medial tympaniform membrane (MTM) and internal labium (LI) was visualized and recorded as the time delay was slowly swept through several oscillation periods. The digitized video frames are shown at increments of 0.2 ms in the time delay. In this case the oscillation frequency was ~ 700 Hz. Note that both the MTM and LI exhibit large motion amplitude. Also labelled are external labium (LE) and bronchi (B).

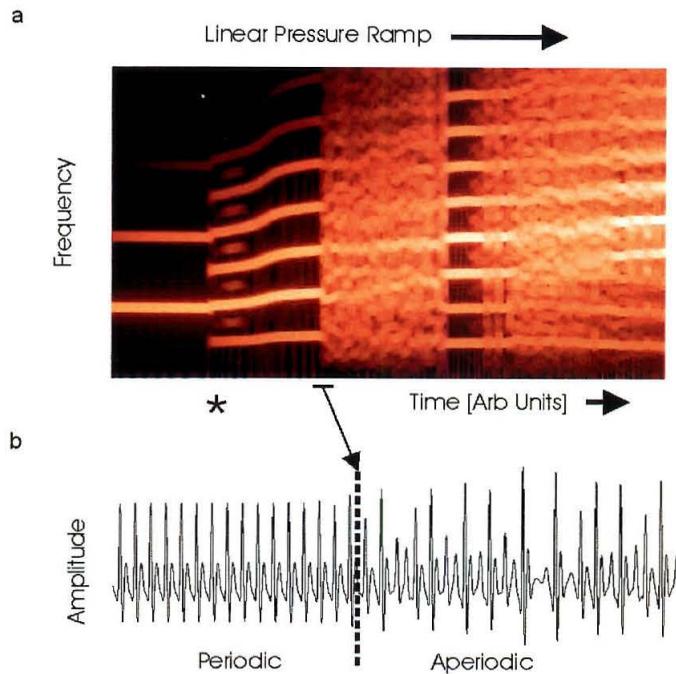


Figure B.4: Results of a simple numerical model of airflow through a compliant constriction. The syringeal membranes are modeled as a system of coupled masses and springs. a) Spectrogram of the oscillatory membrane displacement. The parameter associated with bronchial pressure is ramped slowly and linearly up from zero. Note the onset of oscillation at a threshold pressure, the period doubling events (*), and the rapid transition to 'chaotic' dynamics. b) The membrane displacement shown expanded around the latter transition showing that the transition occurs rapidly, i.e., in roughly one oscillatory period.

part from syringeal dynamics¹, as opposed to vocal tract filtering effects.

We have demonstrated that the isolated syrinx behaves as classic nonlinear dynamical system, often exhibiting rapid transitions between distinct oscillatory states. The strongly nonlinear behavior of the zebra finch syrinx results from the large amplitude motion of the vibrating membranes and the involvement of the relatively heavy labium internum. As a consequence, simple trajectories in a small number of control parameters can result in a complex sequence of acoustic signals. Thus, in addition to the central song control system, the syrinx must also be viewed as an important component in the hierarchy of structures determining the temporal organization of birdsong. Furthermore, the resulting complexity of the mapping from articulators to acoustic output suggests that song learning may require a sophisticated internal model of syringeal dynamics.

Appendix C: List of abbreviations

DPSS Discrete prolate spheroidal sequence.

fMRI Functional magnetic resonance imaging.

ISI Interspike interval.

LFP Local field potential.

LE External labium.

LI Labium internum.

LIP Lateral intra-parietal area.

MEG Magnetoencephalography.

MT Middle temporal area.

MTM Medial tympaniform membrane.

PC Principal component.

PRR Parietal reach region.

PSTH Peri-stimulus time histogram.

SNR Signal-to-noise ratio.

STA Spike-triggered average.

SVD Singular value decomposition.

Bibliography

- H. Akaike. A new look at the statistical model identification. *IEEE Trans. Auto. Cont.*, 19:716–722, 1974.
- D.J Amit. The hebbian paradigm reintegrated: local reverberations as internal representation. *Behav. Brain Sci.*, 18:617–626, 1995.
- R.A. Andersen. Encoding of intention and spatial location in the posterior parietal cortex. *Cerebral Cortex*, 5(5):457–469, 1995.
- A. Baddeley. Working memory. *Science*, 255:556–559, 1992.
- W. Bair, C. Koch, W. Newsome, and K Britten. Power spectrum analysis of bursting cells in area mt in the behaving monkey. *J. Neurosci.*, 14(5):2870–2892, 1994.
- S. Barash, R.M. Bracewell, L. Fogassi, J.W. Gnadt, and R.A. Andersen. Saccade-related activity in the lateral intraparietal area .1. temporal properties - comparison with area 7a. *J. Neurophysiol.*, 66:1095–1108, 1991.
- A.P. Batista and R.A. Andersen. The parietal reach region codes the next planned movement in a sequential reach task. *J. Neurophys.*, 85(2):539–544, 2001.
- A.P. Batista, C.A. Buneo, L.H. Snyder, and R.A. Andersen. Reach plans in eye-centered coordinates. *Science*, 285(5425):257–260, 1999.
- D.A. Berry, H. Hertz, I.R. Titze, and K. Krischer. Interpretation of biomechanical simulations of normal and chaotic vocal fold oscillation with empirical eigenfunctions. *J. Acoust. Soc. Am.*, 95:3595–3604, 1994.
- C. D. Bertram and T.J. Pedley. A mathematical model of unsteady collapsible tube behaviour. *J. Biomechanics*, 15:39–50, 1982.

- G.G. Blasdel and G. Salama. Voltagesensitive dyes reveal a modular organization in monkey striate cortex. *Nature*, 321:579–585, 1986.
- G.J. Blatt, R.A. Andersen, and G.R. Stoner. Visual receptive-field organization and cortico-cortical connections of the lateral intraparietal area (area lip) in the macaque. *J. Comp. Neurol.*, 299:421–445, 1990.
- J.J. Bouyer, M.F. Montaron, and A. Rougeul. *Electroenceph, Clin. Neurophysiol.*, 51:244–252, 1981.
- R.M. Bracewell, P. Mazzonni, S. Barash, and R.A. Andersen. Motor intention activity in the macaque’s lateral intraparietal area .2. changes of motor plan. *J. Neurophysiol.*, 76:1457–1464, 1996.
- S.L. Bressler, R. Coppola, and R. Nakamura. Episodic multiregional cortical coherence at multiple frequencies during visual task-performance. *Nature*, 366:153–156, 1993.
- D. Brillinger. *Developments in Statistics*, volume 1, pages 33–129. Academic Press Inc., 1978.
- D.R. Brillinger. *Time series data analysis and theory*. Holt, Rinehart and Winston, Inc:New York, 1975.
- K.H. Britten, M.N. Shadlen, W.T. Newsome, and J.A. Movshon. The analysis of visual motion: A comparison of neuronal and psychophysical performance. *J. Neurosci*, 12:4745–4765, 1992.
- C.B. Bruce and M.E. Goldberg. Primate frontal eye fields. i. single neurons discharging before saccades. *J. Neurophysiol.*, 53:603–635, 1985.
- J.C. Brumberg, L.G. Nowak, and D.A. McCormick. Ionic mechanisms underlying repetitive high-frequency burst firing in supragranular cortical neurons. *J. Neurosci*, 20:4829–4843, 2000.

- C.A. Buneo, M.R. Jarvis, A.P. Batista, and R.A. Andersen. Properties of spike train spectra in two parietal reach areas. *J. Neurophysiol*, To be submitted.
- S. Cardoso de Oliveira, A. Thiele, and K-P. Hoffmann. Synchronization of neuronal activity during stimulus expectation in a direction discrimination task. *J. Neurosci.*, 17:9248–9260, 1997.
- M.V. Chafee and P.S. Goldman-Rakic. Matching patterns of activity in primate prefrontal area 8a and parietal area 7ip neurons during a spatial working memory task. *J. Neurophysiol.*, 79(6):2919–2940, 1998.
- W. S. Cleveland. *Visualizing data*. Hobart Press, 1993.
- C.L. Colby and M.E. Goldberg. Space and attention in parietal cortex. *Ann. Rev. Neurosci.*, 22:319–349, 1999.
- B.W. Connors and A.R. Kriegstein. Cellular physiology of the turtle visual cortex - distinctive properties of pyramidal and stellate neurons. *J. Neurosci.*, 6:164–177, 1986.
- C.E. Cosans and P.S. Ulinski. Spatial-organization of axons in turtle visual cortex - intralamellar and interlamellar projections. *J. Comp. Neurol.*, 296:548–558, 1990.
- D.R. Cox and V. Isham. *Point processes*. Chapman and Hall, London, UK, 1980.
- D.R. Cox and P.A.W. Lewis. *The statistical analysis of series of events*. Chapman and Hall:London, 1966.
- D.J. Daley and D. Vere-Jones. *An introduction to the theory of point processes*. Springer-Verlag:New York, 1988.
- H.V. Davila, B.M. Salzberg, L.B. Cohen, and A.S. Waggoner. A large change in axon fluorescence that provides a promising method for measuring membrane potential. *Nature, New. Biol.*, 241:159–160, 1973.

- L. Denby and C. L. Mallows. Singular values of large matrices subject to gaussian perturbations. In *Computing Science and Statistics: Proceedings of the 23rd Symposium on the Interface*, pages 54–57. Interface Foundation, 1991.
- W. Denk, J.H. Strickler, and W.W. Webb. Two photon laser scanning fluorescence microscopy. *Science*, 248:73–76, 1990.
- P.H. Desan. *Forebrain of Reptiles*. Karge, Basil, 1988.
- J.P. Donoghue, J.N. Sanes, N.G. Hatsopoulos, and G. Gaal. Neural discharge and local field potential oscillations in primate motor cortex during voluntary movements. *J. Neurophysiol.*, 79(1):159–173, 1998.
- R. O. Duda and P. E. Hart. *Pattern Classification and Scene Analysis*. Wiley, 1973.
- J.R. Duhamel, C.L. Colby, and M.E. Goldberg. The updating of the representation of visual space in parietal cortex by intended eye-movements. *Science*, 255(5040):90–92, 1992.
- R. Eckhorn, R. Bauer, W. Jordan, M. Brosch, W. Kruse, M. Munk, and H.J. Reitboeck. Coherent oscillations - a mechanism of feature linking in the visual cortex - multiple electrode and correlation analyses in the cat. *Biol. Cybern.*, 60(2):121–130, 1988.
- R. Eckhorn, A. Frien, R. Bauer, T. Woelbern, and H. Kehr. High-frequency (60-90 hz) oscillations in primary visual-cortex of awake monkey. *Neuroreport*, 4(3):243–246, 1993.
- B. Efron and T.J. Tibshirani. *An introduction to the bootstrap*. Chapman and Hall:New York, 1993.
- A.K. Engel, P. Konig, C.M. Gray, and W. Singer. Stimulus-dependent neuronal oscillations in cat visual-cortex - intercolumnar interaction as determined by cross-correlation analysis. *Eur. J. Neurosci.*, 2:588–606, 1990.
- G.B. Ermentrout and N. Kopell. Inhibition-produced patterning in chains of coupled non-linear oscillators. *SIAM J. Math.*, 54:478–507, 1994.

- M.S. Fee, B. Shraiman, B. Pesaran, and P.P. Mitra. The role of nonlinear dynamics of the syrinx in the vocalizations of a songbird. *Nature*, 395:67–71, 1998.
- N.H. Fletcher. Mode locking in non-linearly excited inharmonic musical oscillators. *J. Acoust. Soc. Am.*, 64:1566–1569, 1978.
- N.H. Fletcher. Bird song - a quantitative acoustic model. *J. Theor. Biol.*, 135:455–481, 1988.
- S. Friedman-Hill, P.E. Maldonado, and C.M. Gray. Dynamics of striate cortical activity in the alert macaque: I incidence and stimulus-dependence of gamma-band neuronal oscillations. *Cerebral Cortex*, 10:1105–1116, 2000.
- P. Fries, J.H. Reynolds, A.E. Rorie, and R. Desimone. Modulation of oscillatory neuronal synchronization by selective visual attention. *Science*, 291:1560–1563, 2001.
- S. Funahashi, C.B. Bruce, and P.S. Goldman-Rakic. Mnemonic coding of visual space in the monkey's dorsolateral prefrontal cortex. *J. Neurophysiol.*, 61:331–349, 1989.
- J. Fuster. *Memory in cerebral cortex: An empirical approach to neural networks in the human and nonhuman brain*. MIT Press:Cambridge, MA, 1995.
- J.M. Fuster and J.P. Jervey. Neuronal firing in the inferotemporal cortex of the monkey in a visual memory task. *J. Neurosci.*, 2(3):361–375, 1982.
- F. Gabbiani and C. Koch. *Methods in Neuronal Modeling*, chapter Principles of spike train analysis, pages 313–361. MIT Press, Cambridge, MA, 1998.
- J.W. Gnadt and R.A. Andersen. Memory related motor planning activity in posterior parietal cortex of macaque. *Exp. Brain Res.*, 70:216–220, 1988.
- P.S. Goldman-Rakic. Cellular basis of working memory. *Neuron*, 14(3):477–485, 1995.
- F. Goller and O.N. Larsen. A new mechanism of sound generation in songbirds. *Proc. Nat. Acad. Sci. USA*, 94:14787–14791, 1997.

- E.R. Grannan, D.K. Kleinfeld, and H. Sompolinsky. Stimulus-dependent synchronization of neuronal assemblies. *Neural Comp.*, 5:550–569, 1993.
- C.M. Gray, P. Konig, A.K. Engel, and W. Singer. Oscillatory responses in cat visual-cortex exhibit inter-columnar synchronization which reflects global stimulus properties. *Nature*, 338:334–337, 1989.
- C.M. Gray and D.A. McCormick. Chattering cells: Superficial pyramidal neurons contributing to the generation of synchronous oscillations in visual cortex. *Science*, 274:109–113, 1996.
- C.M. Gray and W. Singer. Stimulus-specific neuronal oscillations in orientation columns of cat visual cortex. *Proc. Natl. Acad. Sci. USA*, 86(5):1698–1702, 1989.
- C.H. Greenewalt. *Bird Song: Acoustics and Physiology*. Smithsonian Institution Press, Washington D.C., 1968.
- A. Grinvald, E. Lieke, R.D. Frostig, C.D. Gilbert, and T.N. Wiesel. Functional architecture of cortex revealed by optical imaging of intrinsic signals. *Nature*, 324:361–364, 1992.
- W.C. Hall and F.F. Ebner. *J. Comp. Neurol.*, 140:101–122, 1970.
- M. Hamalainen, R. Hari, R.J. Ilmoniemi, J. Knuutila, and O.V. Lounasmaa. Magnetoencephalography - theory, instrumentation, and applications to noninvasive studies of the working brain. *Rev. Mod. Phys.*, 65(2):413–497, 1993.
- R.S. Hartley and R.A. Suthers. Airflow and pressure during canary song: direct evidence for mini-breaths. *J. Comp. Phys. A*, 165:15–26, 1989.
- D.O. Hebb. *Organization of Behavior*. Wiley:New York, 1949.
- S.B. Heller and P.S. Ulinski. Morphology of geniculocortical axons in turtles of the genera *pseudemys* and *chrysemys*. *Anat. Embryol.*, 175:505–515, 1987.
- H. Hertz, D.A. Berry, I.R. Titze, and M. Saleh. Analysis of vocal disorders with methods from nonlinear dynamics. *J. Speech Hearing Res.*, 37:1008–1019, 1994.

- J.J. Hopfield. Pattern recognition computation using action potential timing for stimulus presentation. *Nature*, 376:33–36, 1995.
- K. Immelman. *Bird Vocalizations*. Cambridge Univ. Press, Cambridge, England, 1969.
- K. Ishizaka and Flanagan J.L. Synthesis of voiced sounds from a two-mass model of the vocal cords. *Bell Systems Technical Journal*, 51:1233–1268, 1972.
- E.A. Jackson. *Perspectives of Nonlinear Dynamics*. Cambridge Univ. Press, 1992.
- M.R. Jarvis and P.P. Mitra. Sampling properties of the spectrum and coherency of sequences action potentials. *Neural Comput.*, 13:717–749, 2001.
- M. Joliot, L. Ribary, and R. Llinas. Human oscillatory brain activity near 40 hz coexists with cognitive temporal binding. *Proc. Natl. Acad. Sci.*, 91(24):11748–11751, 1994.
- H. Killackey, T. Pellmar, and F.F. Ebner. *Fed. Proceed.*, 31:819, 1972.
- D. Kleinfeld and K.R. Delaney. Distributed representation of vibrissa movement in the upper layers of somatosensory cortex revealed with voltage-sensitive dyes. *J. Comp. Neurol.*, 375: 89–108, 1996.
- D. Kleinfeld, K.R. Delaney, M.S. Fee, J.A. Flores, D.W. Tank, and A. Gelperin. Dynamics of propagating waves in the olfactory network of a terrestrial mollusk - an electrical and optical study. *J. Neurophys.*, 72(3):1402–1419, 1994.
- K.W. Koch and J.M. Fuster. Unit-activity in monkey parietal cortex related to haptic perception and temporary memory. *Exp. Brain Res.*, 76(2):292–306, 1989.
- M. Konishi. Pattern generation in birdsong. *Curr Opin Neurobiol*, 4:827–831, 1994.
- A.K. Kreiter and W. Singer. Stimulus-dependent synchronization of neuronal responses in the visual cortex of the awake macaque monkey. *J. Neurosci.*, 16(7):2381–2396, 1996.
- A.R. Kriegstein and B.W. Connors. Cellular physiology of the turtle visual cortex - synaptic properties and intrinsic circuitry. *J. Neurosci.*, 6:178–191, 1986.

- K.K. Kwong, J.W. Belliveau, D.A. Chesler, I.E. Goldberg, R.M. Weiskoff, B.P. Poncelet, D.N. Kennedy, B.E. Hoppel, M.S. Cohen, and R. Turner. Dynamic magnetic resonance imaging of human brain activity during sensory stimulation. *Proc. Natl. Acad. Sci.*, 89: 5675–5679, 1992.
- T. H. Le and X. P. Hu. Retrospective estimation and correction of physiological artifacts in fmri by direct extraction. *Mag. Res. Med.*, 35(3):290–298, 1996.
- M.A. Lebedev and S.P. Wise. Oscillations in the premotor cortex: single-unit activity from awake, behaving monkeys. *Exp. Brain Res.*, 130:195–215, 2000.
- R. Lorente de No. *Physiology of the nervous system*, chapter Cerebral cortex architecture, intracortical connections, motor projections, pages 291–339. Oxford University Press, Oxford, 1938.
- C.L. Mallows. Linear processes are almost gaussian. *J. Appl. Prob.*, 4:313–329, 1967.
- M. E. Mann and J. Park. Global-scale modes of surface temperature. *J. Geophys. Res. Atmos.*, 99(D12):25819–25833, 1994.
- J.E.W. Mayhew, S. Askew, Y. Zheng, J. Porrill, G.W.M. Westby, P. Redgrave, D.M. Rector, and R.M. Harper. Cerebral vasomotion: 0.1hz oscillation in reflected light imaging of neural activity. *Neuroimage*, 4(3):183–193, 1996.
- P.Z. Mazurskaya. *Z. Evolyut. Biokhim. Fiziol.*, 8:617–624, 1972.
- P. Mazzonni, R.M. Bracewell, S. Barash, and R.A. Andersen. Motor intention activity in the macaque’s lateral intraparietal area .1. dissociation of motor plan from sensory memory. *J. Neurophysiol.*, 76:1439–1456, 1996.
- J.S. McCasland. Neuronal control of bird song production. *J. Neurosci*, 7:23–39, 1987.
- E.K. Miller, C.A. Erickson, and R. Desimone. Neural mechanisms of visual working memory in prefrontal cortex of the macaque. *J. Neurosci.*, 16:5154–5167, 1996.

- P. P. Mitra, S. Ogawa, X. P. Hu, and K. Ugurbil. The nature of spatiotemporal changes in cerebral hemodynamics as manifested in functional magnetic resonance imaging. *Mag. Res. Med.*, 37(4):511–518, 1997.
- P.P. Mitra and B. Pesaran. Analysis of dynamic brain imaging data. *Biophys. J.*, 76: 691–708, 1999.
- U. Mitzdorf. Current source-density method and application in cat cerebral cortex: investigation of evoked potentials and eeg phenomena. *Physiological Rev.*, 65:37–100, 1985.
- Y. Miyashita and H.S. Chang. Neuronal correlate of pictorial short-term memory in the primate temporal cortex. *Nature*, 331:68–70, 1988.
- V.N. Murthy and E.E. Fetz. Oscillatory activity in sensorimotor cortex of awake monkeys: Synchronization of local field potentials and relation to behavior. *J. Neurophysiol.*, 76 (6):3949–3967, 1996a.
- V.N. Murthy and E.E. Fetz. Synchronization of neurons during local field potential oscillations in sensorimotor cortex of awake monkeys. *J. Neurophysiol.*, 76(6):3968–3982, 1996b.
- K. Nakamura, A. Mikami, and K. Kubota. Oscillatory neuronal-activity related to visual short-term-memory in monkey temporal pole. *Neuroreport*, 3(1):117–120, 1992.
- F. Nottebohm, D.B. Kelley, and J.A. Paton. Connections of vocal control nuclei in the canary telencephalon. *J. Comp. Neurol.*, 207, 1982.
- F. Nottebohm, T.M. Stokes, and C.M. Leonard. Central control of song in the canary, *serinus canarius*. *J. Comp. Neurol.*, 165:457–486, 1976.
- S. Ogawa, D.W. Tank, R. Menon, J.M. Ellerman, S.G. Kim, H. Merkle, and K. Ugurbil. Intrinsic signal changes accompanying sensory stimulation - functional mapping with magnetic resonance imaging. *Proc. Natl. Acad. Sci.*, 89(13):5951–5955, 1992.

- D. Orbach, E. V. Obrien, R. Everson, E. Kaplan, B. Knight, and L. Sirovich. Time-course and wavelength dependence of the intrinsic optical signal in visual cortex investigated with a modified karhunen-loeve procedure. *Invest. Opth. Vis. Sci.*, 36(4):S873, 1995.
- H.S. Orbach and L.B. Cohen. Optical monitoring of activity from many areas of the in vitro and in vivo salamander olfactory bulb - a new method for studying functional organization in the vertebrate central nervous system. *J. Neurosci.*, 3:2251–2262, 1983.
- J.B. Pawley. *Handbook of Biological Confocal Miroscopy*. Plenum Press, N.Y. and London, 1995.
- D.B. Percival and A.T. Walden. *Spectral analysis for physical applications*. Cambridge University Press, Cambridge, UK, 1993.
- J.S. Pezaris. *Local circuitry in LIP*. PhD thesis, California Institute of Technology, 2000.
- J.S. Pezaris, M. Sahani, and R.A. Andersen. Tetrodes for monkeys. In J.M. Bower, editor, *Computational Neuroscience*. Plenum Press:New York, 1997.
- J.S. Pezaris, M. Sahani, and R.A. Andersen. Response-locked changes in auto- and cross-covariations in parietal cortex. *Neurocomp.*, 26-27:471–476, 1999.
- J. Prechtl, L. B. Cohen, B. Pesaran, P. P. Mitra, and D. Kleinfeld. Visual stimuli induce waves of electrical activity in turtle cortex. *Proc. Natl. Acad. Sci.*, 94(14):7621–7626, 1997.
- J.C. Prechtl. Visual motion induces synchronous oscillations in turtle visual cortex. *Proc. Natl. Acad. Sci. USA*, 91:12467–12471, 1994.
- J.C. Prechtl and T.H. Bullock. *Proc. 2nd Joint Symp. on Neural Comp., UCSD and Caltech*, 5:105–114, 1995.
- P.H. Price. Developmental determinants of structure in zebra finch song. *J. Comp. and Physiol. Psych.*, 93:260–277, 1979.

- C.R. Rao. *Linear Statistical Inference and its Applications*. Wiley:New York, 1965.
- P. E. Rapp, A. M. Albano, I. D. Zimmerman, and M. Mimenez. Phase-randomized surrogates can produce spurious identifications of nonrandom structure. *Phys. Lett. A*, 192(1):27–33, 1994.
- M.L. Reece and J. O’Keefe. The tetrode: An improved technique for multi-unit extracellular recording. In *Soc. Neurosci. Abs.*, volume 15, 1989.
- P.R. Roelfsema, A.K. Engel, P. Konig, and W. Singer. Visuomotor integration is associated with zero time-lag synchronization among cortical areas. *Nature*, 385(6612):157–161, 1997.
- J.R. Rosenberg, A.M. Amjad, P. Breeze, D.R. Brillinger, and D.M. Halliday. The fourier approach to the indentification of functional coupling between neuronal spike trains. *Prog. Biophys. Molec. Biol.*, 53:1–31, 1989.
- A.L. Roskies. The binding problem. *Neuron*, 24:7–9, 1999.
- A. Rouguel, J.J. Bouyer, L. Dedet, and O. Debray. Fast somato-parietal rhythms during combined focal attention and immobility in baboon and squirrel monkey. *Electro. Clin. Neurophysiol.*, 46:310–319, 1979.
- M. Sahani. *Latent variable models for neural data analysis*. PhD thesis, California Institute of Technology, 1999.
- M. Sahani, J.S. Pezaris, and R.A. Andersen. On the separation of signals from neighboring cells in tetrode recordings. In M.I. Jordan, M.J. Kearns, and S.A. Solla, editors, *Advances in Neural Information Processing Systems 10*. MIT Press:Cambridge, MA, 1998.
- J.N. Sanes and J.P. Donoghue. Oscillations in local-field potentials of the primate motor cortex during voluntary movement. *Proc. Natl. Acad. Sci. USA*, 90:4470–4474, 1993.

- C. Scharff and F. Nottebohm. A comparative study of the behavioral deficits following lesions of various parts of the zebra finch song system: Implications for song learning. *J. Neurosci*, 11:2896–2913, 1991.
- A. Sengupta and P. P. Mitra. Distributions of singular values from some random matrices. *Phys. Rev. E*, 60(3):3389–3392, 1999.
- D.M. Senseman. Correspondence between visually evoked voltage-sensitive dye signals and synaptic activity recorded in cortical pyramidal cells with intracellular microelectrodes. *Vis Neurosci*, 13:963–977, 1996.
- M.N. Shadlen and W.T. Newsome. Noise, neural codes and cortical organization. *Curr. Opin. Neurobiol.*, 4:569–579, 1994.
- K.V. Shenoy, S.A. Kureshi, D. Meeker, B.L. Gillikin, A.P. Batista, C.B. Buneo, S. Cao, J.W. Burdick, and R.A. Andersen. Toward prosthetic systems controlled by parietal cortex. In *Soc. Neurosci. Abs.*, volume 25, 1999.
- W. Singer and C.M. Gray. Visual feature integration and the temporal correlation hypothesis. *Ann. Rev. Neurosci.*, 18:555–586, 1995.
- D. Slepian and H. O. Pollak. Prolate spheroidal wavefunctions fourier analysis and uncertainty i. *Bell Sys. Tech. Journal*, 40:43–63, 1961.
- L.H. Snyder, A.P. Batista, and R.A. Andersen. Coding of intention in the posterior parietal cortex. *Nature*, 386(6621):167–170, 1997.
- L.H. Snyder, A.P. Batista, and R.A. Andersen. Saccade-related activity in the parietal reach region. *J Neurophysiol*, 83(2):1099–1102, 1999.
- H. Sompolinsky and M. Tsodyks. Segmentation by a network of oscillators with stored memories. *Neural Comp.*, pages 642–657, 1994.

- M. Steriade, I. Timofeev, N. Durmuller, and F. Grenier. Dynamic properties of corticothalamic neurons and local circuit interneurons generating fast rhythmic (30–40 Hz) spike bursts. *J. Neurophysiol.*, 79:483–490, 1998.
- C. Tallon-Baudry and O. Bertrand. Oscillatory gamma activity in humans and its role in object representation. *Trends Cog. Sci.*, 3(4):151–162, 1999.
- C. Tallon-Baudry, O. Bertrand, F. Peronnet, and J. Pernier. Induced gamma-band activity during the delay of a visual short-term memory task in humans. *J. Neurosci.*, 18(11):4244–4254, 1998.
- C. Tallon-Baudry, A. Kreiter, and O. Bertrand. Sustained and transient oscillatory responses in the gamma and beta bands in a visual short-term memory task in humans. *Vis. Neurosci.*, 16(3):449–459, 1999.
- J. Theiler and P. E. Rapp. Reexamination of the evidence for low-dimensional, nonlinear structure in the human electroencephalogram. *EEG Clin. Neurophys.*, 98(3):213–222, 1996.
- D. J. Thomson. Time series analysis of holocene climate data. *Phil. Trans. Roy. Soc. Lond. A*, 330:601–616, 1990.
- D. J. Thomson and A. D. Chave. *Advances in Spectrum Analysis and Array Processing*, volume 1, chapter Jackknifed Error Estimates for Spectra, Coherences, and Transfer Functions, pages 58–113. Prentice Hall, 1991.
- D.J. Thomson. Spectrum estimation and harmonic analysis. *Proc. IEEE*, 70:1055–1096, 1982.
- D.J. Thomson. Projection filters for data analysis. In *Proc. Seventh IEEE SP Work. Stat. Sig. and Array Proc.*, pages 39–42, Quebec, Canada, 1994.
- N. Ueda and R. Nakano. Mixture density estimation via em algorithm with deterministic annealing. *Proc. IEEE: Neural networks and signal processing*, 69:69–77, 1994.

- P.S. Ulinski. Organization of corticogeniculate projections in the turtle, *pseudemys scripta*. *J. Comp. Neurol.*, 254:529–542, 1986.
- C. von der Malsberg and C. Schneider. A neural cocktail-party processor. *Biol. Cybern.*, 54:29–40, 1986.
- E. Vu, M.E. Mazurek, and Y.-C. Kuo. Identification of a forebrain motor programming network for the learned song of zebra finches. *J. Neurosci.*, 14:6924–6934, 1994.
- P. D. Welch. The use of fast fourier transform for the estimation of power spectra: A method based on time averaging over short, modified periodograms. *IEEE Trans. Aud. Electro.*, 15:70–73, 1967.
- A. Yu and Margoliash D. Temporal hierarchical control of singing in birds. *Science*, 273: 1871–1875, 1996.
- Y.D. Zhou and J.M. Fuster. Mnemonic neuronal activity in somatosensory cortex. *Proc. Natl. Acad. Sci. USA*, 93(19):10533–10537, 1996.