

Spectral Analysis for Neural Signals

Bijan Pesaran, PhD

Center for Neural Science, New York University
New York, New York

Introduction

This chapter introduces concepts fundamental to spectral analysis and applies spectral analysis to characterize neural signals. Spectral analysis is a form of time series analysis and concerns a series of events or measurements that are ordered in time. The goal of such an analysis is to quantitatively characterize the relationships between events and measurements in a time series. This quantitative characterization is needed to derive statistical tests that determine how time series *differ* from one another and how they are *related* to one another. Time series analysis comprises two main branches: time-domain methods and frequency-domain methods. Spectral analysis is a frequency-domain method for which we will use the multitaper framework (Thomson, 1982; Percival and Walden, 1993). The treatment here also draws on other sources (Brillinger, 1978; Jarvis and Mitra, 2001; Mitra and Bokil, 2007).

In this chapter, we will focus on relationships within and between one and two time series, known as univariate and bivariate time series. Throughout this discussion, we will illustrate the concepts with experimental recordings of spiking activity and local field potential (LFP) activity. The chapter on Multivariate Neural Data Sets will extend the treatment to consider several simultaneously acquired time series that form a multivariate time series, such as in imaging experiments. The chapter on “Application of Spectral Methods to Representative Data Sets in Electrophysiology and Functional Neuroimaging” will review some of this material and present additional examples.

First we begin by motivating a particular problem in neural signal analysis that frames the examples in this chapter. Second, we introduce signal processing and the Fourier transform and discuss practical issues related to signal sampling and the problem of aliasing. Third, we present stochastic processes and their characterization through the method of moments. The moments of a stochastic process can be characterized in both the time domains and frequency domains, and we will discuss the relation between these characterizations. Subsequently, we present the problem of scientific inference, or hypothesis testing, in spectral analysis through the consideration of error bars. We finish by considering an application of spectral analysis involving regression.

Motivation

When a microelectrode is inserted into the brain, the main features that are visible in the extracellular potential it measures are the spikes and the rhythms they ride on. The extracellular potential results from

currents flowing in the extracellular space, which in turn are produced by transmembrane potentials in local populations of neurons. These cellular events can be fast, around 1 ms for the action potentials that appear as spikes, and slow, up to 100 ms, for the synaptic potentials that predominantly give rise to the LFP. How spiking and LFP activity encode the sensory, motor, and cognitive processes that guide behavior, and how these signals are related, are fundamental, open questions in neuroscience (Steriade, 2001; Buzsaki, 2006). In this chapter, we will illustrate these analysis techniques using recordings of spiking and LFP activity in macaque parietal cortex during the performance of a delayed look-and-reach movement to a peripheral target (Pesaran et al., 2002). This example should not be taken to limit the scope of potential applications, and other presentations will motivate other examples.

The Basics of Signal Processing

Spiking and LFP activity are two different kinds of time series, and all neural signals fall into one of these two classes. LFP activity is a continuous process and consists of a series of continuously varying voltages in time, x . Spiking activity is a point process, and, assuming that all the spike events are identical, consists of a sequence of spike times. The counting process, N_t , is the total number of events that occur in the process up to a time, t . The mean rate of the process, λ , is given by the number of spike events divided by the duration of the interval. If we consider a sufficiently short time interval, $\delta t = 1$ ms, either a spike event occurs or it does not. Therefore, we can represent a point process as the time derivative of the counting process, dN_t , which gives a sequence of delta functions at the precise time of each spike, t_n . We can also represent the process as a sequence of times between spikes, $\tau_n = t_{n+1} - t_n$, which is called an interval process. These representations are equivalent but capture different aspects of the spiking process. We will focus on the counting process, dN_t . $dN_t = 1 - \lambda \delta t$ when there is a spike and $dN_t = -\lambda \delta t$ elsewhere. Note that these expressions correct for the mean firing rate of the process. As we will see, despite the differences between point and continuous processes, spectral analysis treats them in a unified way. The following sections present some basic notions that underlie statistical signal processing and time series analysis.

Fourier transforms

Time series can be represented by decomposing them into a sum of elementary signals. One domain for signals, which we call the time domain, is simply each point in time. The time series is represented by its amplitude at each time point. Another domain is the frequency domain and consists of sinusoidal

NOTES

functions, one for each frequency. The process is represented by its amplitude and phase at each frequency. The time and frequency domains are equivalent, and we can transform signals between them using the Fourier transform. Fourier transforming a signal that is in the time domain, x_t , will give the values of the signal in the frequency domain, $\tilde{x}(f)$. The tilde denotes a complex number with amplitude and phase.

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

Inverse Fourier transforming $\tilde{x}(f)$ transforms it to the time domain. To preserve all the features in the process, these transforms need to be carried out over an infinite time interval. However, this is never realized in practice. Performing Fourier transforms on finite duration data segments distorts features in the signal and, as we explain below, spectral estimation employs data tapers to limit these distortions.

Nyquist frequency, sampling theorem, and aliasing

Both point and continuous processes can be represented in the frequency domain. When we sample a process, by considering a sufficiently short interval in time, δt , and measuring the voltage or the presence or absence of a spike event, we are making an assumption about the highest frequency in the process. For continuous processes, the sampling theorem states that when we sample an analog signal that is band-limited, so that it contains no frequencies greater than the Nyquist rate (B Hz), we can perfectly reconstruct the original signal if we sample at a sampling rate, $F_s = \frac{1}{\delta t}$, of at least $2B$ Hz. The

original signal is said to be band-limited because it contains no energy outside the frequency band given by the Nyquist rate. Similarly, once we sample a signal at a certain sampling rate, F_s , the maximum frequency we can reconstruct from the sampled

signal is called the Nyquist frequency, $\frac{1}{2} F_s$.

The Nyquist frequency is a central property of all sampled, continuous processes. It is possible to sample the signal more frequently than the bandwidth of the original signal, with a sampling rate greater than twice the Nyquist rate, without any problems. This is called oversampling. However, if we sample the signal at less than twice the Nyquist rate, we cannot reconstruct the original signal without errors. Errors arise because components of the signal that exist at a higher frequency than $\frac{1}{2} F_s$ become aliased into

lower frequencies by the process of sampling the signal. Importantly, once the signal has been sampled at F_s , we can no longer distinguish between continuous processes that have frequency components greater

than the Nyquist frequency of $\frac{1}{2} F_s$. To avoid the

problem of aliasing signals from high frequencies into lower frequencies by digital sampling, an anti-aliasing filter is often applied to continuous analog signals before sampling. Anti-aliasing filters act to low-pass the signal at a frequency less than the Nyquist frequency of the sampling.

Sampling point processes does not lead to the same problems as sampling continuous processes. The main consideration for point processes is that the sampled point process be orderly. Orderliness is achieved by choosing a sufficiently short time interval so that each sampling interval has no more than one event. Since this means that there are only two possible outcomes to consider at each time step, 0 and 1, analyzing an orderly point process is simpler than analyzing a process that has multiple possible outcomes, such as 0, 1, and 2, at each time step.

Method of moments for stochastic processes

Neural signals are variable and stochastic owing to noise and the intrinsic properties of neural firing. Stochastic processes (also called random processes) can be contrasted with deterministic processes, which are perfectly predictable. Deterministic processes evolve in exactly the same way from a particular point. In contrast, stochastic processes are described by probability distributions that govern how they evolve in time. Stochastic processes evolve to give a different outcome, even if all the samples of the process originally started from the same point. This is akin to rolling dice on each trial to determine the neural activity. Each roll of the dice is a realization from a particular probability distribution, and it is this distribution that determines the properties of the signal. When we measure neural signals to repeated trials in an experiment, we assume that the signals we record on each trial are different realizations or outcomes of the same underlying stochastic process.

Another powerful simplification is to assume that the properties of the stochastic process generating the neural signals within each trial are stationary and that their statistical properties don't change with time—even within a trial. This is clearly not strictly true for neural signals because of the nonstationarities in behavior. However, under many circumstances, a reasonable procedure is to weaken the stationarity

assumption to short-time stationarity. Short-time stationarity assumes that the properties of the stochastic process have stationarity for short time intervals, say 300-400 ms, but change on longer time scales. In general, the window for spectral analysis is chosen to be as short as possible to remain consistent with the spectral structure of the data; this window is then 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 $\Delta t = N$ by $\Delta f = W$, then $NW \geq 1$. We can then estimate the statistical properties of the stochastic process by analyzing short segments of data and, if necessary and reasonable, averaging the results across many repetitions or trials. Examples of time-frequency characterizations are given below. Note that this presentation uses “normalized” units. This means that we assume the sampling rate to be 1 and the Nyquist frequency interval to range from $-\frac{1}{2}$ to $\frac{1}{2}$. The chapter Application of Spectral Methods to Representative Data Sets in Electrophysiology and Functional Neuroimaging presents the relationships below in units of time and frequency.

Spectral analysis depends on another assumption: that the stochastic process which generates the neural signals has a spectral representation.

$$x_t = \int_{-\frac{1}{2}}^{\frac{1}{2}} \tilde{x}(f) \exp(2\pi i f t) df$$

Remarkably, the same spectral representation can be assumed for both continuous processes (like LFP activity) and point processes (like spiking activity), so the Fourier transform of the spike train, t_n , is as follows:

$$d\tilde{N}(f) = \sum_{n=1}^N \exp(2\pi i f t_n)$$

The spectral representation assumes that underlying stochastic processes generating the data exist in the frequency domain, but that we observe their realizations as neural signals, in the time domain. As a result, we need to characterize the statistical properties of these signals in the frequency domain: This is the goal of spectral analysis.

The method of moments characterizes the statistical properties of a stochastic process by estimating the moments of the probability distribution. The first moment is the mean; the second moments are the

variance and covariance (for more than one time series), and so on. If a stochastic process is a Gaussian process, the mean and variance or covariances completely specify it. For the spectral representation, we are interested in the second moments. The spectrum is the variance of the following process:

$$S_x(f) \delta(f - f') = E[\tilde{x}^*(f) \tilde{x}(f')]$$

$$S_{dN}(f) \delta(f - f') = E[d\tilde{N}^*(f) d\tilde{N}(f')]$$

The delta function indicates that the process is stationary in time. The asterisk denotes complex conjugation. The cross-spectrum is the covariance of two processes:

$$S_{XY}(f) \delta(f - f') = E[\tilde{x}^*(f) \tilde{y}(f')]$$

The coherence is the correlation coefficient between each process at each frequency and is simply the covariance of the processes normalized by their variances.

$$C_{XY}(f) = \frac{S_{XY}(f)}{\sqrt{S_x(f) S_y(f)}}$$

This formula represents the cross-spectrum between the two processes, divided by the square root of the spectrum of each process. We have written this expression for two continuous processes; analogous expressions can be written for pairs of point-continuous processes or point-point processes by substituting the appropriate spectral representation. Also, we should note that the assumption of stationarity applies only to the time interval during which we carry out the expectation.

Multitaper spectral estimation

The simplest estimate of the spectrum, called the *periodogram*, is proportional to the square of the data sequence, $|\tilde{x}_n(f)|^2$. This spectral estimate suffers from two problems. The first is the problem of bias. This estimate does not equal the true value of the spectrum unless the data length is infinite. Bias arises because signals at different frequencies are mixed together and “blurred.” This bias comes in two forms: *narrow-band bias* refers to bias in the estimate due to mixing signals at different nearby frequencies. *Broad-band bias* refers to mixing signals at different frequencies at distant frequencies. The second problem is the problem of variance. Even if the data length were

NOTES

infinite, the periodogram spectral estimate would simply square the data without averaging. As a result, it would never converge to the correct value and would remain inconsistent.

Recordings of neural signals are often sufficiently limited so that bias and variance can present major limitations in the analysis. Bias can be reduced, however, by multiplying the data by a data taper, w_t , before transforming to the frequency-domain, as follows:

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

Using data tapers reduces the influence of distant frequencies at the expense of blurring the spectrum over nearby frequencies. The result is an increase in narrow-band bias and a reduction in broad-band bias. This practice is justified under the assumption that the true spectrum is locally constant and approximately the same for nearby frequencies. Variance is usually addressed by averaging overlapping segments of the time series. Repetitions of the experiment also give rise to an ensemble over which the expectation can be taken, but this precludes the assessment of single-trial estimates.

An elegant approach toward the solution of both the above problems has been offered by the multitaper spectral estimation method, in which the data are multiplied by not one, but several, orthogonal tapers and Fourier-transformed in order to obtain the basic quantity for further spectral analysis. The simplest example of the method is given by the direct multitaper estimate, $S_{MT}(f)$, defined as the average of individual tapered spectral estimates,

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

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

The $w_t(k)$ ($k = 1, 2, \dots, K$) constitute K orthogonal taper functions with appropriate properties. A particular choice for these taper functions, with optimal spectral concentration properties, is given by the discrete prolate spheroidal sequences, which we will call “Slepian functions” (Slepian and Pollack, 1961). Let $w_t(k, W, N)$ be the k th Slepian function of length N and frequency bandwidth parameter W . The Slepian functions would then form an orthogonal basis set for sequences of length, N , and be 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 having 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 the formula

$$U(f) = \sum_{t=1}^N w_t \exp(-2\pi i f t). \text{ Then we can consider}$$

the problem of finding sequences w_t so that the spectral amplitude $U(f)$ is maximally concentrated in the interval $[-W, W]$. Maximizing this concentration parameter, subject to constraints, yields a matrix eigenvalue equation for $w_t(k, W, N)$. The eigenvectors of this equation are the Slepian functions. The remarkable fact is that the first $2NW$ eigenvalues $\lambda_k(N, W)$ (sorted in descending order) are each approximately equal to 1, while the remainder are approximately zero. The Slepian functions can be shifted in concentration from $[-W, W]$ centered around zero frequency to any nonzero 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 to take the leading $2NW - 1$ Slepian functions as data tapers in the multitaper analysis. The remaining functions have progressively worsening spectral concentration properties. For illustration, in the left column of Figure 1, we show the first four Slepian functions for $W = 5/N$. In the right column, we show the time series example from the earlier subsection multiplied by each of the successive data tapers. In the left column of Figure 2, we show the spectra of the data tapers themselves, displaying the spectral concentration property. The vertical marker denotes the bandwidth parameter W . Figure 2 also shows the magnitude-squared Fourier transforms of the tapered time series presented in Figure 1. The arithmetic average of these spectra for $k = 1, 2, \dots, 9$ (note that only 4 of 9 are shown in Figs. 1 and 2) gives a direct multitaper estimate of the underlying process.

Figure 3A shows the periodogram estimate of the spectrum based on a single trial of LFP activity during the delayed look-and-reach task. The variability in the estimate is significant. Figure 3B presents the multitaper estimate of the spectrum on the same data with $W = 10$ Hz, averaged across 9 tapers. This

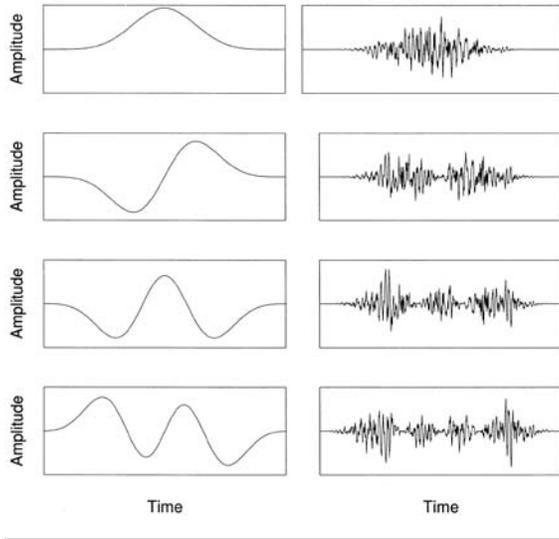


Figure 1. Slepian functions in the time domain. Left panels: First four Slepian functions for $NW = 5$. Right panels: Data sequence multiplied by each Slepian data taper on left.

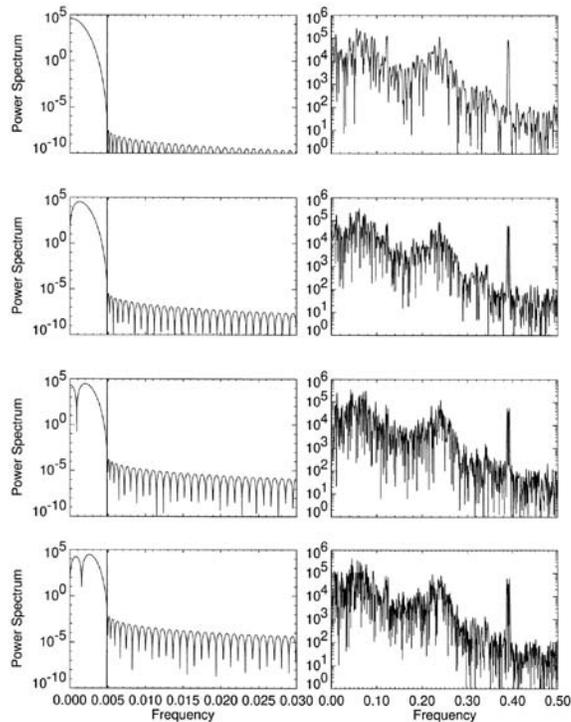


Figure 2. Slepian functions in the frequency domain. Left panel: spectra of Slepian functions from left panels of Figure 1. Right panel: spectra of data from right panels of Figure 1.

estimate is much smoother and reveals the presence of two broad peaks in the spectrum, at 20 Hz and 60 Hz. Figure 3C shows the multitaper spectrum estimate on the same data with $W = 20$ Hz. This estimate is even smoother than the 10 Hz, which reflects the increased number of tapers available to

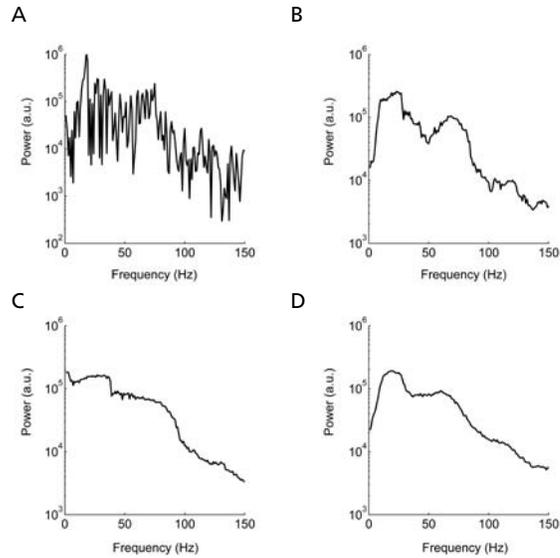


Figure 3. Spectrum of LFP activity in macaque lateral intraparietal area (LIP) during delay period before a saccade-and-reach to preferred direction. **A**, Single trial, 500 ms periodogram spectrum estimate. **B**, Single trial, 500 ms, 10 Hz multitaper spectrum estimate. **C**, Single trial, 500 ms, 20 Hz multitaper spectrum estimate. **D**, Nine-trial average, 500 ms, 10 Hz multitaper spectrum estimate. a.u. = arbitrary units.

average across (19 tapers instead of 9). However, the assumption that the spectrum is constant with the 20 Hz bandwidth is clearly wrong and leads to noticeable distortion in the spectrum, in the form of a narrow-band bias. Figure 3D shows the multitaper estimate with bandwidth $W = 10$ Hz averaged across 9 trials. Compared with Figure 3B, this estimate is noticeably smoother and contains the same frequency resolution. This series illustrates the advantages of multitaper estimates and how they can be used to improve spectral resolution.

Bandwidth selection

The choice of the time window length N and the bandwidth parameter W is critical for applications. No simple procedure can be given for these choices, which in fact depend on the data set at hand, and are best made iteratively using visual inspection and some degree of trial and error. $2NW$ 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 $2NW$. 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 by increasing narrow-band bias. This can be done formally by trading off an appropriate weighted sum of the estimated variance and bias. However, as a rule, we find fixing the time bandwidth product NW at a small number (typically

NOTES

3 or 4), and then varying the window length in time until sufficient spectral resolution is obtained, to be a reasonable strategy. It presupposes that the data are examined in the time-frequency plane so that N may be significantly smaller than the total data length.

Figure 4 illustrates these issues using two spectrogram estimates of the example LFP activity averaged across 9 trials. Each trial lasts approximately 3 s and consists of a 1 s baseline period, followed by a 1–1.5 s delay period, during which a movement is being planned. The look-reach movement is then executed. Each spectrogram is shown with time on the horizontal axis, frequency on the vertical axis, and power color-coded on a log base-10 scale. Figure 4A shows the spectrogram estimated using a 0.5 s duration analysis window and a 10 Hz bandwidth. The time-frequency tile this represents is shown in the white rectangle. This estimate clearly shows the sustained activity following the presentation of the spatial cue at 0 s that extends through the movement's execution. Figure 4B shows a spectrogram of the same data estimated using a 0.2 s duration analysis window and a 25 Hz bandwidth. The time-frequency tile for this estimate has the same area as Figure 4A, so each estimate has the same number of degrees of freedom. However, there is great variation in the time-frequency resolution trade-off between these estimates: Figure 4B better captures the transients in the signal, at the loss of significant frequency resolution that distorts the final estimate. Ultimately, the best choice of time-frequency resolution will depend on the frequency band of interest, the temporal dynamics in the signal, and the number of trials available for increasing the degrees of freedom of a given estimate.

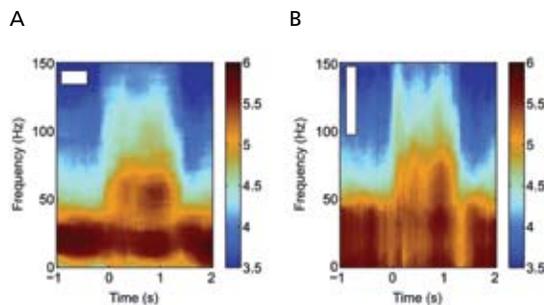


Figure 4. Spectrogram of LFP activity in macaque LIP averaged across 9 trials of a delayed saccade-and-reach task. Each trial is aligned to cue presentation, which occurs at 0 s. Saccade and reach are made at around 1.2 s. **A**, Multitaper estimate with duration of 500 ms and bandwidth of 10 Hz. **B**, Multitaper estimate with duration of 200 ms and bandwidth 25 Hz. White rectangle shows then time-frequency resolution of each spectrogram. The color bar shows the spectral power on a log scale in arbitrary units.

Calculating error bars

The multitaper method confers one important advantage: 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 within a time series. Error bars can be constructed using a number of procedures, but broadly speaking, there are two types. The fundamental notion common to both types of error bars is the local frequency ensemble. That is, if the spectrum of the process is locally flat over a bandwidth $2W$, then the tapered Fourier transforms $\tilde{x}_k(f)$ constitute a statistical ensemble for the Fourier transform of the process at the frequency, f_o . This locally flat assumption and the orthogonality of the data tapers mean that the $\tilde{x}_k(f)$ are uncorrelated random variables having the same variance. 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 first type of error bar is the asymptotic error bar. For large N , $\tilde{x}_k(f)$ may be assumed to be asymptotically, normally distributed under some general circumstances. As a result, the estimate of the spectrum is asymptotically distributed according to a χ^2_{dof}

distribution scaled by $\frac{S(f)}{dof}$. The number of degrees of freedom (dof) is given by the total number of data tapers averaged to estimate the spectrum. This would equal the number of trials multiplied by the number of tapers.

For the second type of error bar, we can use the local frequency ensemble to estimate jackknife error bars for the spectra and all other spectral quantities (Thomson and Chave, 1991; Wasserman, 2007). The idea of the jackknife is to create different estimates by, in turn, leaving out a data taper. This creates a set of spectral estimates that forms an empirical distribution. A variety of error bars can be constructed based on such a distribution. If we use a variance-stabilizing transformation, the empirical distribution can be well approximated using a Gaussian distribution. We can then calculate error bars according to the normal interval by estimating the variance of the distribution and determining critical values that set the error bars. Inverting the variance-stabilizing transformation gives us the error bars for the original spectral estimate. This is a standard tool in statistics and provides a more conservative error bar than the asymptotic error bar. Note that the degree to which the two error bars agree constitutes a test of how well the empirical distribution follows the asymptotic distribution. The variance-stabilizing transforma-

tion for the spectrum is the logarithm. The variance-stabilizing transformation for the coherence, the magnitude of the coherency, is the arc-tanh.

As an example, Figure 5A shows asymptotic and Figure 5B empirical jackknife estimates of the spectral estimate illustrated in Figure 3D. These are 95% confidence intervals and are largely the same between the two estimates. This similarity indicates that, for these data, the sampling distribution of the spectral estimate follows the asymptotic distribution across trials and data tapers. If we were to reduce the estimate's number of degrees of freedom by reducing the number of trials or data tapers, we might expect to see more deviations between the two estimates, with the empirical error bars being larger than the asymptotic error bars.

Correlation functions

Neural signals are often characterized in terms of correlation functions. Correlation functions are equivalent to computing spectral quantities but with important statistical differences. For stationary processes, local error bars can be imposed for spectral estimates in the frequency domain. This is not true for correlation functions, even assuming stationarity, because error bars for temporal correlation functions are nonlocal. Nonlocality in the error bars means that uncertainty about the correlation function at one lag is influenced by the value of the correlation function across other lags. The precise nature of the nonlocality relies on the temporal dependence within the underlying process. Consequently, correlation function error bars must be constructed by assuming there are no dependencies between different time bins. This is a far more restrictive assumption than the one holding that neighboring frequencies are locally flat and rarely achieved in practice. Other problems associated with the use of correlation functions are

that if the data contain oscillatory components, they are compactly represented in frequency space and lead to nonlocal effects in the correlation function. Similar arguments apply to the computation of correlation functions for point and continuous processes. One exception is for spiking examples in which there are sharp features in the time-domain correlation functions, e.g., owing to monosynaptic connections.

Figure 6 illustrates the difference between using spectral estimates and correlation functions. Figure 6A shows the spectrum of spiking activity recorded in macaque parietal cortex during a delay period before a coordinated look-and-reach. The duration of the spectral estimate is 500 ms, the bandwidth is 30 Hz, and the activity is averaged over nine trials. Thin lines show the empirical 95% confidence intervals. Figure 6B shows the auto-correlation function for the same data, revealing some structure around short lags and inhibition at longer lags. There is a hint of some ripples, but the variability in the estimate is too large to see them clearly. This is not too surprising, because the correlation function estimate is analogous to the periodogram spectral estimate, which also suffers from excess statistical variability. In contrast, the spectrum estimate clearly reveals the presence of significant spectral suppression and a broad spectral peak at 80 Hz. The dotted line shows the expected spectrum from a Poisson process having the same rate.

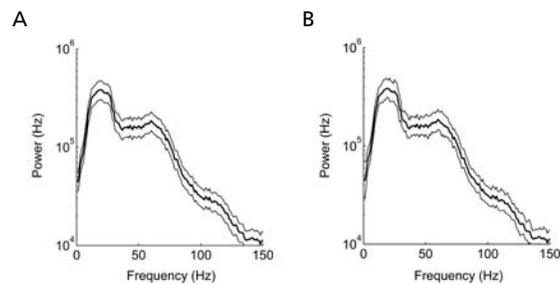


Figure 5. 95% confidence error bars for LFP spectrum shown in Figure 3D. **A**, Asymptotic error bars assuming chi-squared distribution. **B**, Empirical error bars using leave-one-out jackknife procedure.

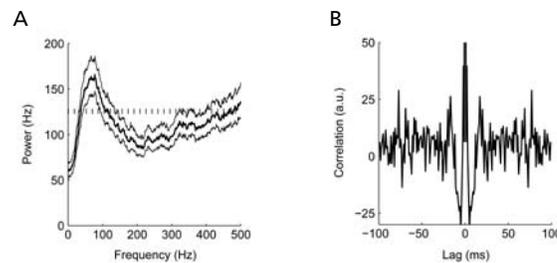


Figure 6. Spike spectrum and correlation function of spiking activity in macaque LIP during delay period before a saccade-and-reach to the preferred direction. **A**, Multitaper spectrum estimate duration of 500 ms; bandwidth 15 Hz averaged across 9 trials. Thin lines show 95% confidence empirical error bars using leave-one-out procedure. Dotted horizontal line shows firing rate. **B**, Correlation function estimate from which shift predictor has been subtracted. a.u. = arbitrary units.

Coherence

The idea of a local frequency ensemble motivates multitaper estimates of the coherence between two-point or continuous processes. Given two time series

NOTES

and the corresponding multiple tapered Fourier transforms $\tilde{x}_k(f)$, $\tilde{y}_k(f)$, the following direct estimates can be defined for the coherence function:

$$C_{XY}(f) = \frac{\frac{1}{K} \sum_k \tilde{x}_k^*(f) \tilde{y}_k(f)}{\sqrt{S_x(f) S_y(f)}}$$

This definition allows us to estimate the coherence from a single trial. Estimating the coherence presents many of the same issues as estimating the spectrum, except that more degrees of freedom are needed to ensure a reasonable estimate. In common with spectrum estimates, the duration and bandwidth of the estimator need to be chosen to allow sufficient degrees of freedom in the estimator. Increasing the number of trials will increase the effective resolution of the estimate.

Figure 7 shows the coherence and correlations between two simultaneously recorded spike trains from macaque parietal cortex averaged over nine trials. Figure 7A shows the coherence estimated with 16 Hz bandwidth. The horizontal dotted line represents expected coherence for this estimator when there is no coherence between the spike trains. The coherence significantly exceeds this threshold, as shown by the 95% confidence intervals, in a broad frequency band. Figure 7B illustrates the coherence estimated with a 30 Hz bandwidth. The variability in the estimate is reduced, as is the noise floor of the estimator, as shown by the lower horizontal dotted line. Figure 7C shows the cross-correlation function for these data. Here, too, there

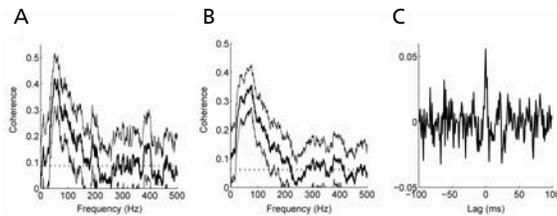


Figure 7. Spike coherence and cross-correlation function for spiking activity of two simultaneously recorded neurons in macaque LIP during delay period before a saccade-and-reach to the preferred direction for both cells. **A**, Multitaper coherence estimate duration of 500 ms; bandwidth 16 Hz averaged across 9 trials. Thin lines show 95% confidence empirical error bars using leave-one-out procedure. Dotted horizontal line shows expected coherence under the null hypothesis that coherence is zero. **B**, Multitaper coherence estimate duration of 500 ms; bandwidth 30 Hz averaged across 9 trials. Conventions as for **A**. **C**, Cross-correlation function estimate from which shift predictor has been subtracted.

is structure in the estimate, but the degree of variability lowers the power of the analysis.

Regression using spectral feature vectors

Detection of period signals is an important problem that occurs frequently in the analysis of neural data. Such signals can arise as a result of periodic stimulation and can manifest as 50/60 Hz line noise. We pursue the effects of periodic stimulation in the multivariate case in the next chapter Multivariate Neural Data Sets: Image Time Series, Allen Brain Atlas. As discussed therein, certain experiments that have no innate periodicity may also be cast into a form that makes them amenable to analysis as periodic stimuli. We now discuss how such components may be detected and modeled in the univariate time series by performing a regression on the spectral coefficients.

Periodic components are visible in preliminary estimates as sharp peaks in the spectrum, which, for multitaper estimation with Slepian, appear with flat tops owing to narrow-band bias. Consider one such sinusoid embedded in colored noise:

$$x(t) = A \cos(2\pi f t + \phi) + \eta(t)$$

It is customary to apply a least-squares procedure to obtain A and ϕ , by minimizing the sum of squares

$\sum_t |x(t) - A \cos(2\pi f_0 t + \phi)|^2$. However, this is a nonlinear procedure that must be performed numerically; moreover, it effectively assumes a white-noise spectrum. Thomson's F-test offers an attractive alternative within the multitaper framework by reducing the line-fitting procedure to a simple linear regression.

Starting with a data sequence containing N samples, multiplying both sides of the equation by a Slepian taper $w_k(t)$ with bandwidth parameter $2W$, and Fourier transforming, one obtains this result:

$$\tilde{x}_k(f) = \mu U_k(f - f_0) + \mu^* U_k(f + f_0) + N_k(f)$$

Here $\mu = A \exp(i\phi)$ and $U_k(f)$ and $N_k(f)$ are the Fourier transforms of $w_k(f)$ and $\eta(t)$, respectively. If f_0 is larger than the bandwidth W , then $f - f_0$ and $f + f_0$ are separated by more than $2W$, and $U_k(f - f_0)$ and $U_k(f + f_0)$ have minimal overlap. In that case, one can set $f = f_0$ and neglect the $U_k(f + f_0)$ term to obtain the following linear regression equation at $f = f_0$:

$$\tilde{x}_k(f) = \mu U_k(0) + N_k(f_0)$$

The solution is given by

$$\hat{\mu}(f_0) = \frac{\sum_{k=1}^K U_k(0) \tilde{x}_k(f_0)}{\sum_{k=1}^K |U_k(0)|^2}$$

The goodness of fit of this model may be tested using an F-ratio statistic with $(2, 2K - 2)$ degrees of freedom, which is usually plotted as a function of frequency to determine the position of the significant sinusoidal peaks in the spectrum,

$$F(f) = \frac{(K-1) |\hat{\mu}(f)|^2 \sum_{k=1}^K |U_k(0)|^2}{\sum_{k=1}^K |\tilde{x}_k(f) - \hat{\mu}(f) U_k(0)|^2}$$

Once an F-ratio has been calculated, peaks deemed significant by the F-test, and which exceed the significance level $1-1/N$, may be removed from the original process in order to obtain a reshaped estimate of the smooth part of the spectrum:

$$S_{\text{reshaped}}(f) = \frac{1}{K} \sum_{k=1}^K \left| \tilde{x}_k(f) - \sum_i \mu_i U_k(f - f_i) \right|^2$$

This reshaped estimate may be augmented with the previously determined line components, to obtain a so-called mixed spectral estimate. This provides one of the more powerful applications of the multi-taper methodology, since, in general, estimating such mixed spectra is difficult.

References

- Brillinger DR (1978) Comparative aspects of the study of ordinary time series and of point processes. In: *Developments in statistics* vol. 11. Orlando, FL: Academic Press.
- Buzsaki G (2006) *Rhythms of the Brain*. New York: Oxford UP.
- Jarvis MR, Mitra PP (2001) Sampling properties of the spectrum and coherency of sequences of action potentials. *Neural Comput* 13:717-749.
- Mitra PP, Bokil H (2007) *Observed Brain Dynamics*. New York: Oxford UP.
- Percival DB, Walden AT (1993) *Spectral analysis for physical applications*. Cambridge, UK: Cambridge UP.

Pesaran B, Pezaris JS, Sahani M, Mitra PP, Andersen RA (2002) Temporal structure in neuronal activity during working memory in macaque parietal cortex. *Nat Neurosci* 5:805-811.

Slepian D, Pollack HO (1961) Prolate spheroidal wavefunctions: Fourier analysis and uncertainty I. *Bell System Tech Journal* 40:43-63.

Steriade M (2001) *The intact and sliced brain*. Cambridge, MA: MIT Press.

Thomson DJ (1982) Spectrum estimation and harmonic analysis. *Proc IEEE* 70:1055-1096.

Thomson DJ, Chave AD (1991) Jackknifed error estimates for spectra, coherences, and transfer functions. In: *Advances in spectrum analysis and array processing*, pp 58-113. Englewood Cliffs, NJ: Prentice Hall.

Wasserman L (2007) *All of nonparametric statistics*. New York: Springer-Verlag.