

Variability of Extracellular Spike Waveforms of Cortical Neurons

MICHALE S. FEE, PARTHA P. MITRA, AND DAVID KLEINFELD

Bell Laboratories, Lucent Technologies, Murray Hill, New Jersey 07974; and Department of Physics, University of California, La Jolla, California 92093

SUMMARY AND CONCLUSIONS

1. Here we study the variability in extracellular records of action potentials. Our work is motivated, in part, by the need to construct effective algorithms to classify single-unit waveforms from multiunit recordings.

2. We used microwire electrode pairs (stereotrodes) to record from primary somatosensory cortex of awake, behaving rat. Our data consist of continuous records of extracellular activity and segmented records of extracellular spikes. Spectral and principal component techniques are used to analyze mean single-unit waveforms, the variability between different instances of a single-unit waveform, and the underlying background activity.

3. The spectrum of the variability between different instances of a single-unit waveforms is not white, and falls off above 1 kHz with a frequency dependence of roughly f^{-2} . This spectrum is different from that of the mean spike waveforms, which falls off roughly as f^{-4} , but is essentially identical with the spectrum of background activity. The spatial coherence of the variability on the 10- μm scale also falls off at high frequencies.

4. The variability between different instances of a single-unit waveform is dominated by a relatively small number of principal components. As a consequence, there is a large anisotropy in the cluster of the spike waveforms.

5. The background noise cannot be represented as a stationary Gaussian random process. In particular, we observed that the spectrum changes significantly between successive 20-ms intervals. Furthermore, the total power in the background activity exhibits larger fluctuations than is consistent with a stationary Gaussian random process.

6. Roughly half of the single-unit spike waveforms exhibit systematic changes as a function of the interspike interval. Although this results in a non-Gaussian distribution in the space of waveforms, the distribution can be modeled by a scalar function of the interspike interval.

7. We use a set of 44 mean single-unit waveforms to define the space of differences between spike waveforms. This characterization, together with that of the background activity, is used to construct a filter that optimizes the detection of differences between single-unit waveforms. Further, an information theoretic measure is defined that characterizes the detectability.

INTRODUCTION

Much of the study of neuronal activity relies on the inference of the spiking output from individual neurons on the basis of measurements of their extracellular signals. However, extracellular recordings of brain activity often contain signals from more than one neuron. Because neighboring neurons often have quite different physiological properties, it is usually desirable to discriminate the signal from one or more individual neurons that contribute to the signal. This discrimination is based on differences in the details of the extracellular action potential waveforms as a consequence

of the type and spatial distribution of currents in the cell and as a function of the position and geometry of the electrode. These differences provide a means to classify different waveforms as belonging to the same neuron. However, extracellular sources of noise and intrinsic spike-to-spike variability obscure the classification process.

A priori, we expect that there are at least two sources of variability that may contribute to the observed shape of the extracellular waveform of a given neuron. The first originates in extracellular currents from other cells. Every part of a neuron, e.g., axons, dendrites, and synapses, is capable of generating currents (Llinas 1988), so there are many possible microscopic sources of noise. If these sources are uncorrelated with the observed spike, at least on the millisecond time scale, this contribution may be viewed as additive noise. A second source of variability may be associated with systematic changes in the spike waveform. For example, changes in the height and width of the action potential have been observed in successive spikes of a burst, as seen in some layer 5 pyramidal neurons in vitro (McCormick et al. 1985).

Here we analyze extracellular records from layers 2/3 through layer 6 of primary somatosensory vibrissa cortex in rat. We ask the following questions. 1) What is the spectral composition of the spike waveform variability and the background neuronal activity? 2) Is the spectral composition stationary across time? 3) Is the amplitude distribution of the variability Gaussian? In particular, systematic variation in the shape of the spike waveform may lead to an apparent non-Gaussian distribution. 4) Can we use these results to construct an optimal filter for the detection of differences between waveforms from different single units?

Our motivation is twofold. On the one hand, we suggest that systematic patterns of variability in the extracellular signal may be useful for the in vivo classification of neuronal type. On the other hand, the decomposition of multiunit extracellular signals into contributions from individual units is fundamentally dependent on the statistics of waveform variability both extrinsic and intrinsic to the neuron. We suggest that these statistics have not been properly accounted for in past work.

Preliminary accounts of this work have appeared (Mitra et al. 1995).

METHODS

Electrophysiology

We recorded regular- and fast-spiking units (Simons 1978) from layers 2 through 6 of the vibrissal area of primary somatosensory cortex of Long-Evans rats. Four independently adjustable stereo-

electrodes (McNaughton et al. 1983) were implanted through the intact dura mater into neocortex. In brief, the electrodes were constructed from a twisted pair of 25-mm polyamide-coated tungsten wires (California Fine Wire, Grover City, CA). The ends of the electrodes were cut with sharp scissors at a 45° angle and gold-plated. Electrode impedances in physiological saline were typically 0.1 MΩ at 1 kHz for both reactive and resistive components, and electrodes produced a Johnson noise of ~30 nV√Hz over the frequency band of interest. Signals were buffered near the head of the animal with field effect transistors (NB Labs, Denville, TX), amplified (×10⁴), band-pass filtered between 300 Hz (5-pole Bessel high-pass filter) and 10 kHz (8-pole constant-phase low-pass filter; Frequency Devices, Haverhill, MA), and digitized at 25 kHz with a 12-bit digital-to-analog converter (no. DT2821; Data Translation, Marlboro, MA) that had an effective resolution of 10 bits. The acquisition was controlled by the "Discovery" data acquisition program (Datawave Technologies, Longmont, CO). The care and experimental manipulation of our animals were in strict accord with guidelines from the National Institutes of Health (1985) and have been reviewed and approved by the Institutional Animal Care and Use Committee at Bell Laboratories.

Data acquisition

Data were acquired in either of two modes: continuous acquisition mode or segmented acquisition mode. In continuous mode, the digitized voltage signals from stereotrode pairs are continuously recorded to disk. The signals for each pair are denoted $V_x(t)$ and $V_y(t)$, where x and y label the wire and t is a discrete variable; an example from a particular pair is shown in Fig. 1a. Sections of data containing spikes were selected with the use of a threshold crossing criterion. We extracted a segment of 64 samples with the peak of the spike centered at sample 11 (see Fig. 1a, vertical lines); each segment defines a pair of vectors denoted $\mathbf{V}_x^{(n)} = \{V_x(t + \tau_n)\}_{t=-10}^{T-10}$ and $\mathbf{V}_y^{(n)} = \{V_y(t + \tau_n)\}_{t=-10}^{T-10}$, where τ_n is the time of the peak of the n th instance of the waveform and $T = 64$. Spike waveforms were sorted on the basis of the amplitude at the peak of the waveform on each wire, i.e., peak $\{\mathbf{V}_x^{(n)}\}$ versus peak $\{\mathbf{V}_y^{(n)}\}$ (Fig. 1a, inset). A small sample of the sorted waveforms is shown in Fig. 1b and the autocorrelation function of the arrival times for the entire set of sorted waveforms is shown in Fig. 1c. The autocorrelation shows a clear suppression of spikes at short intervals, consistent with the spike train of a single unit. For the purposes of our analysis of waveform variability, we include only extracellular records with one or two well-isolated single units.

The difference between a particular instance of a spike waveform and the mean waveform is defined as a spike residual. We calculate the residuals for all sorted waveforms as follows. 1) The segmented spike waveforms (see above) are resampled to place the center of mass of their peak at sample 11 (Fee et al. 1996) and the mean waveform is calculated as the average of the centered, segmented waveforms. This procedure removes the dominant source of jitter in computing the average waveform. 2) The mean waveform is resampled by cubic spline interpolation to generate a template with 0.8-μs resolution. Temporally shifted versions of the mean waveform are generated by shifting the template in 0.8-μs steps and resampling at the 40-μs sample period. We thus generate a set of 50 mean waveforms, each shifted in time by 0.8 μs, that span the 40-μs sample period. 3) Each of the shifted means is subtracted from the spike waveform, and the residual with the minimum total squared error is kept; a sample of spike residuals and their amplitude distribution is shown in Fig. 1, d and e, respectively. In addition, a segment of 64 samples of the continuous record in the interval between 4.0 and 1.44 ms before the onset of each spike waveform is extracted for the purpose of analyzing background activity.

In segmented acquisition mode, a threshold crossing of either signal of the stereotrode pair triggers the acquisition of the spike waveforms on both wires, as described in Fee et al. (1996). In contrast to the case of continuous acquisition, only 32 samples of the waveform from each

of the two wires of the electrode may be saved. The voltage sample with the largest amplitude (waveform peak) is set as the fifth sample, the data are recentered as described in Fee et al. (1996), and the companion waveform is shifted in register; a time stamp saved with each waveform indicates the time of the waveform peak with 100-μs resolution. Segmented acquisition was used to obtain the relatively large number of waveforms required for our results on optimal filtering; in this case as many as four single units per wire were isolated, as described in Fee et al. (1996).

Spectral analysis

We used the direct multitaper estimation techniques of Thomson (1982) to calculate the power spectral density, denoted $S_{xx}(f)$, and the coherence, denoted $R_{xy}(f)$, from the measured voltage signals. These spectral measures are defined by

$$S_{xx}(f) = \langle |\tilde{\mathbf{V}}_x|^2 \rangle \quad (1)$$

and

$$R_{xy}(f) = \frac{\langle \tilde{\mathbf{V}}_x \tilde{\mathbf{V}}_y^* \rangle}{\sqrt{\langle |\tilde{\mathbf{V}}_x|^2 \rangle \langle |\tilde{\mathbf{V}}_y|^2 \rangle}} \quad (2)$$

where $|\tilde{\mathbf{V}}|^2 \equiv \tilde{\mathbf{V}}\tilde{\mathbf{V}}^*$ and $\langle \dots \rangle$ denotes an average over all tapers, i.e.

$$\langle \tilde{\mathbf{V}}\tilde{\mathbf{U}} \rangle = \frac{1}{N} \frac{1}{K} \sum_{n=1}^N \sum_{k=1}^K \tilde{\mathbf{V}}^{(n,k)} \tilde{\mathbf{U}}^{(n,k)} \quad (3)$$

where $\tilde{\mathbf{V}}^{(n,k)} = \{\tilde{V}^{(n,k)}(f)\}_{f=0}^{f_N}$ is the discrete Fourier transform of $\mathbf{V}^{(n)}$ multiplied by the k th window function, or taper, $\mathbf{W}_k = \{w_k(t)\}_{t=-10}^{T-10}$ (see below), i.e., $\tilde{V}^{(n,k)}(f) = \sum_{t=-10}^{T-10} \exp[i2\pi ft] w_k(t) V^{(n)}(t)$, N is the number of instances of the waveform (~10³ to 10⁴ in the present work), and K is the number of tapers (2 or 3 in the present work). The Nyquist frequency is $f_N = (2t_s)^{-1}$ where t_s is the time per sample (40 μs in the present work). For the spike residuals, the above formula holds with $\mathbf{V}^{(n)}$ replaced with $\delta\mathbf{V}^{(n)}$, where $\delta\mathbf{V}^{(n)} = \mathbf{V}^{(n)} - \langle \mathbf{V} \rangle$.

The use of multiple tapers yields K independent estimates of the spectrum, which are averaged to form a final spectrum (Eq. 3). The frequency resolution of the spectrum, defined in terms of half-width of the spectral bands, Δf , satisfies

$$(2 \cdot \Delta f) \cdot (T \cdot t_s) = K + 1 \quad (4)$$

For spike waveforms acquired in continuous mode ($T = 64$), $\Delta f \sim 600$ Hz. Note that additional smoothing, but no change in bandwidth, is obtained by averaging the spectra from multiple instances (Eq. 3).

The multitaper methods offer advantages (Percival and Walden 1993) that are particularly critical to the estimation of spectra for the mean spike waveforms. In particular 1) the spectra have a large dynamic range of amplitudes. The sequences used to construct the tapers, w_k , are an orthogonal set of functions (discrete prolate spheroidal sequences) that minimize the leakage of power between frequency bands. 2) The segmented spike waveforms contain a peak that, by construction, occurs at one end of the record. The use of multiple tapers, even two, provides a relatively balanced weight across all regions of a record, as opposed to the preferential weight given to the center of a record with only a single taper.

Principal components

The set of waveform vectors $\mathbf{V}^{(n)}$ exists in a T -dimensional space. We consider the directions, known as the principal components, that minimize the covariance between the projections of vectors (Ahmed and Rao 1975; Golub and Van Loan 1989). The correlation or covariance matrix is denoted \mathbf{C} , where

$$\mathbf{C} = \langle \mathbf{V}\mathbf{V}^T \rangle \quad (5)$$

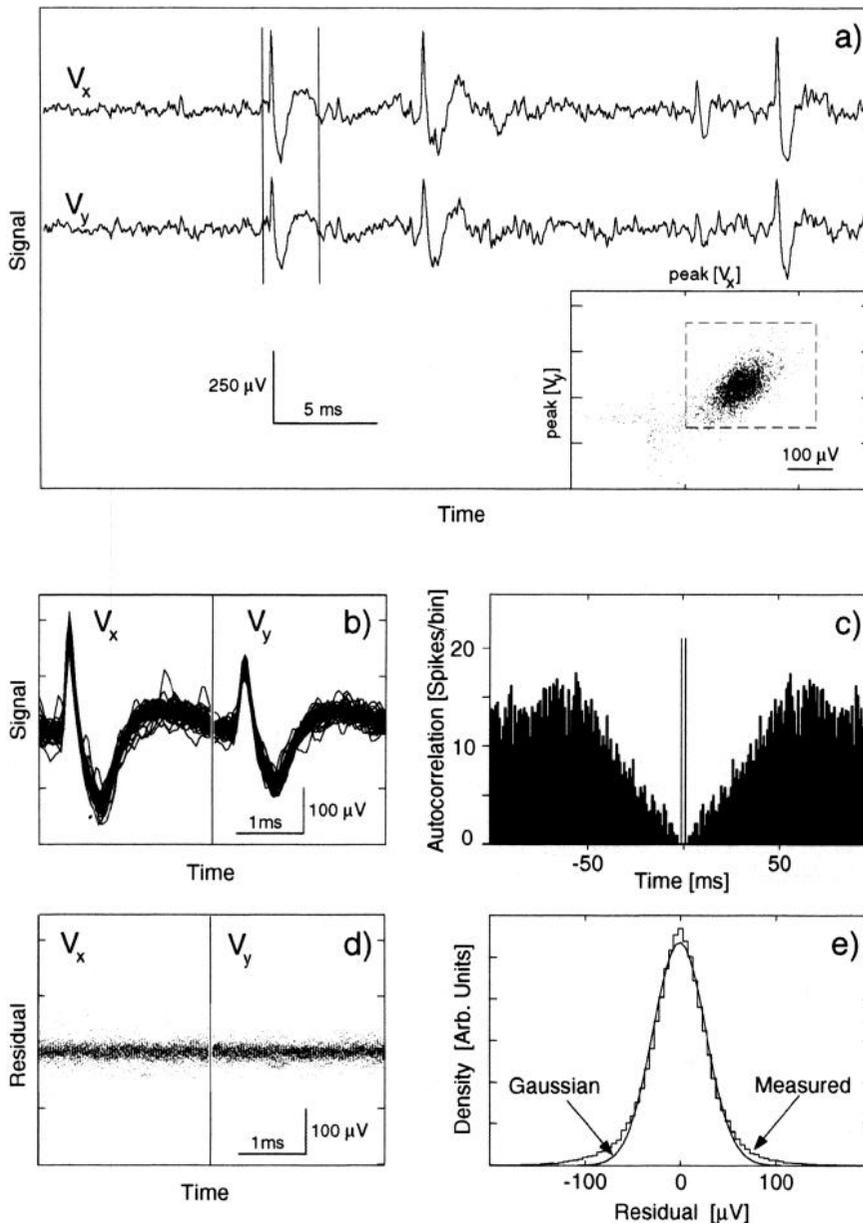


FIG. 1. Representative stereotrode data. *a*: voltage signals from each of the 2 microwires, denoted V_x and V_y , as a function of time. Vertical lines segment a region that corresponds to a spike waveform. *Inset*: distribution of spike waveforms as a function of the maximum value of $V_x^{(n)}(t)$ and $V_y^{(n)}(t)$, for all instances of the waveforms ($N = 2,600$). Only those waveforms within the dashed box are kept. Note that the distribution of points in the projection for this figure appears artificially broad. *b*: sample of 100 sorted spike waveforms from the data in *a*. Note the high degree of overlap among all waveforms. *c*: autocorrelation function of the spike arrival times for all instances of the waveforms. The absence of any amplitude at equal times implies that all of the waveforms could originate from the same single unit. *d*: distribution of residuals as a function of time for each microwire. These residuals are calculated by subtracting the mean waveform from each instance of the spike; 50 waveforms were used for this calculation. *e*: distribution of residuals for all sample points shown in *d*; note that it is nearly Gaussian.

is a real symmetrical T by T matrix with elements $C(t, t') = \langle V(t)V(t') \rangle$; the average is over all instances and satisfies $N > T$. The principal components, denoted U_α with $\alpha = 1, \dots, T$, are the eigenvectors of C , i.e.

$$CU_\alpha = \lambda_\alpha^2 U_\alpha \quad (6)$$

where λ_α^2 is the α th eigenvalue (variance); the order of the eigenvectors is chosen so that the magnitude of the eigenvalues decrease monotonically. From a different perspective, the U_α s are the columns of a unitary matrix $U = [U_1, U_2, \dots, U_T]$ that diagonalizes C , so that

$$U^T C U = \Lambda^2 \quad (7)$$

where Λ^2 is a diagonal matrix with diagonal elements $\lambda_1^2, \lambda_2^2, \dots, \lambda_T^2$. Last, the T -dimensional vectors $V^{(n)}$ and $U^T V^{(n)}$ form a transform pair; the former is the representation of the waveform in time and the latter is the representation in the space of principal compo-

nents. Identical considerations hold with $V^{(n)}$ replaced by $\delta V^{(n)}$, etc.

For the special case in which the waveforms are translationally invariant in time, so that $C(t, t') = C(t - t')$, the diagonal elements of the correlation matrix form the Fourier transform of the power spectrum, truncated to T points in time, and the eigenvectors can be calculated solely from the spectral properties of the waveforms. For the further limit of an infinite length time series, i.e., $T \rightarrow \infty$, the principal components are sinusoids and the above transforms (Eq. 6 and 7) reduce to Fourier transforms.

RESULTS

Spectral properties

We consider first the spectral density (Eq. 1) of the spike waveforms and associated variability, and focus on 2,600 spike waveforms recorded on one wire for the regular-spike-

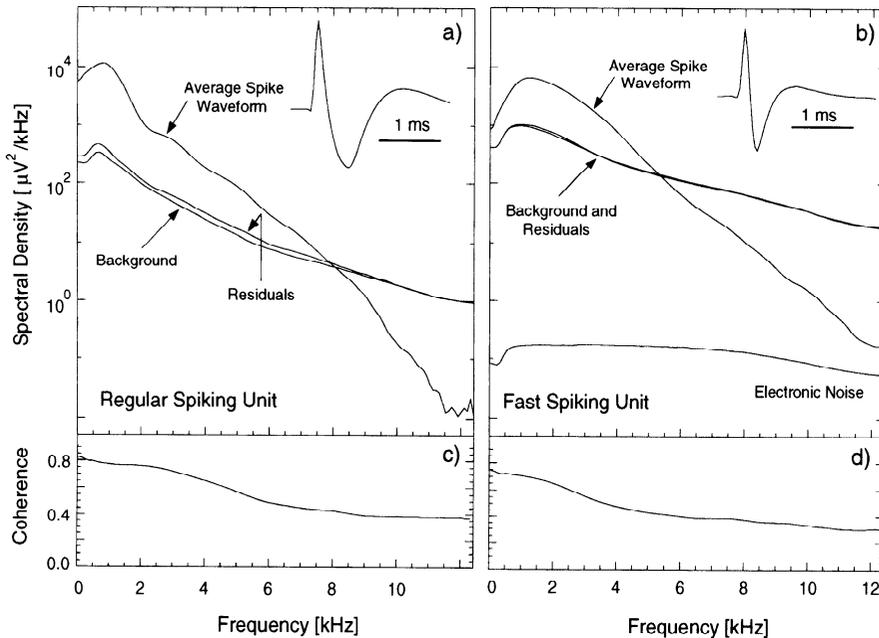


FIG. 2. Spectral analysis of spike waveforms and variability. *a*: power spectral density for a regular-spiking unit (*inset*), $S_{xx}(f)$. Note the close correspondence between the power in the spike residuals and that in background activity taken from the same continuous record. *b*: power spectral density for a fast-spiking unit (*inset*). The frequency dependence of the electronic noise is essentially the frequency response of the electronics. *c* and *d*: spectral coherence, $R_{xy}(f)$, across the 2 wires of the stereotrode pair for the units in *a* and *b*, respectively.

ing unit presented in Fig. 1. We observed that the spectral density of the mean spike waveform has a broad maximum at low frequencies ($f \sim 500$ Hz) and falls off monotonically above ~ 1 kHz roughly as $S_{xx}(f) \propto f^{-4}$ (Fig. 2*a*). The average spectrum of spike residuals has a smaller amplitude at low frequencies and falls off with a frequency dependence of $S_{xx}(f) \propto f^{-2}$. The average spectral density of the background activity is essentially identical to that of the spike residuals. Last, the spectrum of the filtered electronic noise is lower in amplitude than any of our signals (Fig. 2*b*). The electrical noise is white over the frequency range in question (METHODS), so that the frequency dependence of the filtered spectrum is determined by the filters; the spectral properties of the spike residuals are thus minimally affected by the electronic filters.

The spectral density observed for the average waveform of a fast-spiking unit is shown in Fig. 2*b*. As in the case of the regular-spiking unit, the spectral density falls off monotonically at frequencies above ~ 2 kHz, although with a steeper frequency dependence than that observed for regular-spiking units. Further, the spike residuals and background activity have similar spectral properties to those shown in Fig. 2*a*. In general, the above results (Fig. 2, *a* and *b*) are characteristic of those found for all regular-spiking units (10) and fast-spiking units (5) for which we performed spectral analysis.

Clues to the origin of waveform variability come from the coherence (Eq. 2) of the signals across the stereotrode pair. A source of variability common to both wires of the stereotrode would lead to a coherence of 1, whereas an independent noise source on each wire would result in a measured coherence close to 0, or more precisely $\sim 10^{-2}$ for our system.¹ The two signals from the stereotrodes show a

high degree of correlation at frequencies < 3 kHz, i.e., $R_{xy}(f) \sim 0.8$, which falls to $R_{xy}(f) < 0.4$ at higher frequencies (Fig. 2, *c* and *d*). This implies that the high-frequency sources of variability are more localized on the scale of the separation between the stereotrode wires, $25 \mu\text{m}$, than are the low-frequency sources.

The nonwhite spectrum of the waveform residuals implies a degree of correlation between the individual samples of the residuals. The structure of these correlations is revealed by examining the principal components of the covariance matrix of the spike residuals. The eigenvalues and eigenvectors were calculated (Eqs. 5 and 6) for the covariance matrix of the 2,600 spike residuals for the single unit used in the spectral analysis of Fig. 2*a* (Fig. 3*a*). The histogram of sorted eigenvalues is seen to fall off rapidly as a function of component number, with three to four dimensions accounting for half of the total variance (Fig. 3*a*, ●). Further, the variance along the largest principal component of the spike residuals is typically > 3 orders of magnitude larger than the variance along the smallest principal component. In contrast to the spectrum for spike waveforms, the sorted spectrum for isotropically distributed Gaussian noise is relatively flat (P. P. Mitra and A. M. Sengupta, unpublished result), with roughly 20 dimensions accounting for half of the total variance and a ratio of largest to smallest principal component of ~ 2 (Fig. 3*a*). All single-unit clusters we have observed yield a similarly anisotropic eigenvalue spectrum of the covariance matrix. These results imply that the distribution of spike waveform residuals is highly anisotropic in the 64-dimensional space of the residuals.

The eigenvectors associated with the first five principal components of the spike residual waveforms for are shown in Fig. 3*b*. The components are dominated by frequencies that are low compared with the Nyquist frequency, $(2t_s)^{-1} = 12.5$ kHz, which indicates the presence of a large degree of temporal correlation between the waveform residuals at different sample times. Last, the correlation matrix for the spike waveforms is not translationally invariant in time be-

¹ The limiting coherence is given by the inverse of the square root of the number of independent spectral estimates. The latter number is equal to the number of tapers, K , times the number of samples, N , so that for random data one would have $R_{xy}(f) \leq {}^{-1}\sqrt{KM} = {}^{-1}\sqrt{2 \cdot 2,600} \cong 1.4 \cdot 10^{-2}$.

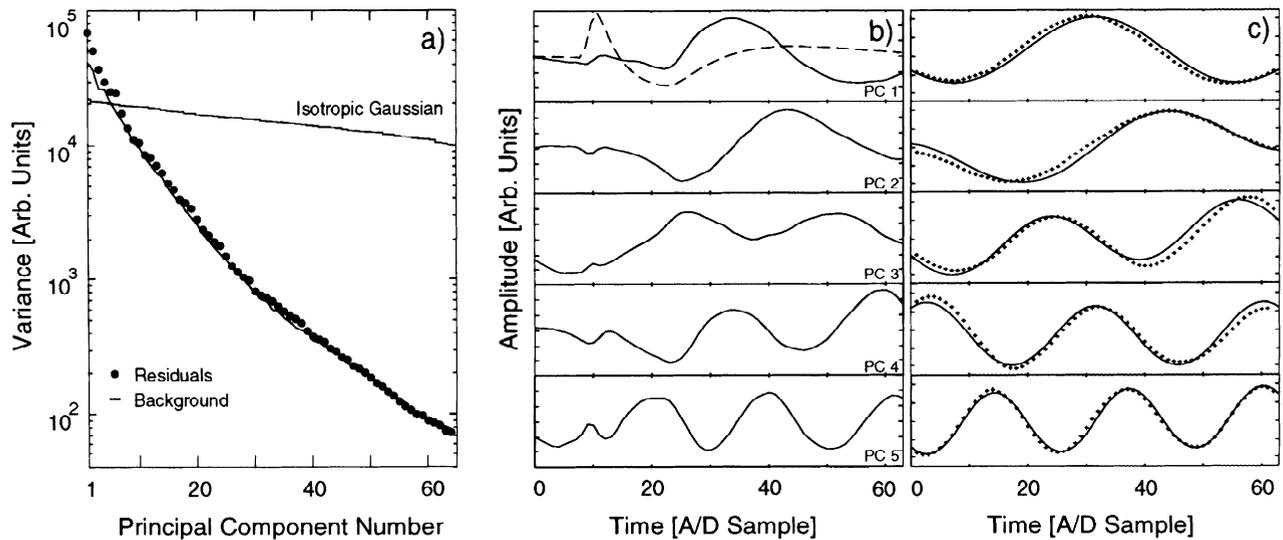


FIG. 3. Principal component analysis of spike waveform residuals and background activity. *a*: eigenvalue spectrum for the spike residuals (●) and the background activity (—), sorted by amplitude. The nearly horizontal line is the spectrum for white noise, normalized so its integral is equal to that for the residuals. *b*: 1st 5 eigenmodes or principal components of the spike residuals. Note the mixture of frequencies in all modes. Dashed line: mean spike waveform. *c*: 1st 5 eigenmodes for the background activity. Dotted lines: eigenmodes calculated from the power spectrum; the close correspondence between the 2 shows that the response is stationary.

cause some variability is locked in time to the spike waveform. Thus the eigenvectors are not consistent with those of a stationary Gaussian signal (Fig. 3*b*).

The spectrum of the background activity, like that of the spike waveform residuals, is also nonwhite (Fig. 2, *a* and *b*). As in the case of the spike waveforms, the sorted eigenvalue spectrum falls off rapidly as a function of component number; in fact, the two spectra are essential identical for the higher principal components (Fig. 3*a*). On the other hand, the segmented waveforms are expected to be translationally invariant for the case of background activity in that there is no relation between the temporal location of excised regions and features in the data record. For this case the eigenvectors can be calculated solely from the power spectrum (METHODS). We observed a good match between the eigenvectors calculated directly from the correlation matrix of the background activity and those calculated indirectly from the power spectrum of the activity (Fig. 3*c*). This comparison acts as a self-consistency check on our numerical methods and demonstrates that the background activity is not locked to our sample window.

Time dependence

Up to now we have considered the average properties of the spike waveform and the variability. We now examine the time dependence of the background activity to determine whether or not the variability can be modeled as a stationary Gaussian process. Three consecutive 20-ms segments of the voltage signal on one wire of the stereotrode during an epoch in which no spikes are observed are shown in Fig. 4*a*. Qualitatively, the signal appears quite different between segments; this difference is highlighted through a comparison of the spectra for each segment with each other and with the average spectrum for 4,000 such segments (Fig. 4*b*). In particular, the spectra for the 20-ms segments have substantial peaks in the subkilohertz frequency range, whereas the average

spectrum is smooth. In general, significant spectral peaks occur in most of the 4,000 segments recorded.

As a second measure of nonstationarity, we consider changes in total power, as opposed to changes in spectral content, as a function of time. The total power in consecutive 20-ms segments, calculated from the integral of the individual power spectra, i.e., $S_{xx} = \sum_f S_{xx}(f)$, is shown in Fig. 4*c* for a 4-s epoch. We compared the fluctuations in the integrated power with those expected for a stationary random process with the same average power spectrum;² the integrated power for an epoch of the derived stationary signal is shown in Fig. 4*c* (gray line). The fluctuations in the total power of the stationary signal are clearly smaller than those in the measured signal. To quantify this difference, the distribution of the total power in each segment of the measured signal and the derived stationary signal were calculated. The power distribution of the measured signal is roughly 3 times as wide at the 10 and 90% points as is the distribution of the stationary Gaussian random process (Fig. 4*d*), with a significant tail at high power. In toto, we find that the background activity cannot be modeled as a stationary Gaussian random process. To the extent that the background activity is the dominant contribution to the variability among spike waveforms, this variability is not stationary.

Systematic variability

The distribution of projections of the spike residuals onto the principal components of the residuals need not be

² The fluctuations for the equivalent stationary process were calculated as follows: the spectra for 4,000 segments of numerically generated Gaussian random noise were constructed, each consisting of 500 samples (20 ms ÷ 40 μs per sample). The spectrum for each segment was computed and then multiplied by the average spectrum determined for the background activity (Fig. 4*b*). The distribution of power in this set of derived spectra was then calculated (Fig. 4*d*).

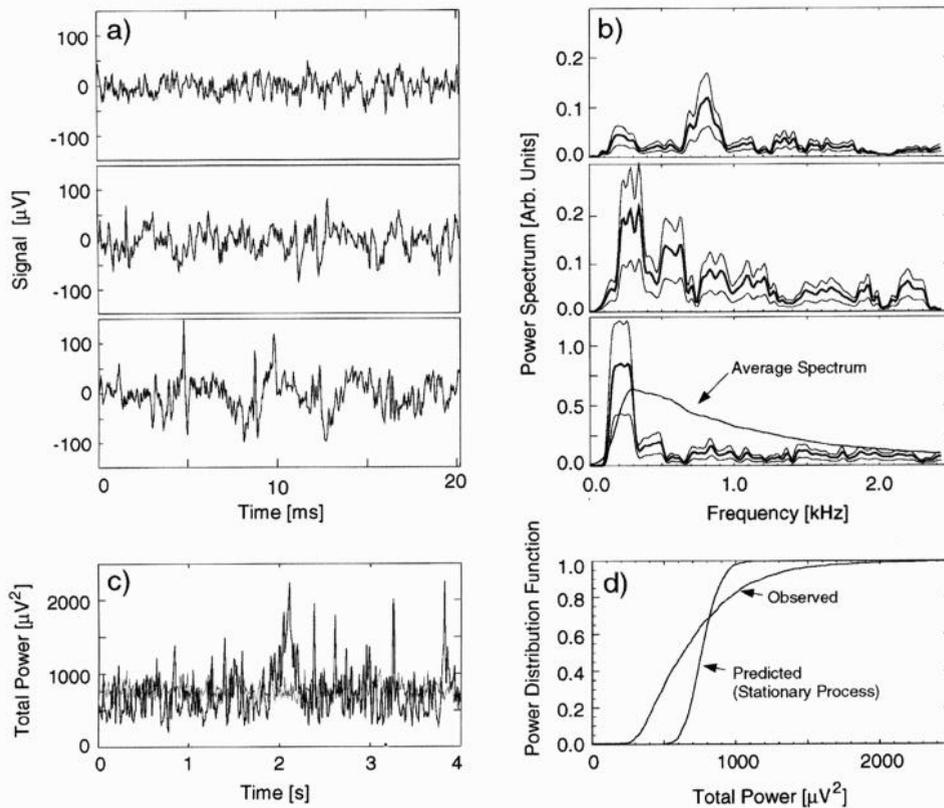


FIG. 4. Nonstationarity of background activity. *a*: voltage waveforms for 3 20-ms epochs of background activity. *b*: power spectra of the epochs shown in *a*. *c*: integrated power for successive 20-ms epochs. Gray line: power for an equivalent stationary Gaussian random process (see text for details). *d*: distribution function for the measured power and the derived power for an equivalent stationary process. Note the greater spread in amplitudes for the observed spectrum.

Gaussian, although the distribution of residuals at each time sample may be nearly Gaussian, as seen for the data (Fig. 1*e*). In particular, we observed that the distributions of projections of the residuals onto the first few principal components of the variability are skewed for about half of the units we examined. For the higher principal components, the distribution of projections is generally not significantly different from Gaussian.

The apparent non-Gaussian distribution of the dominant modes may be the consequence of systematic changes in the spike waveform over time. One mechanism for change is a slow drift in the position of the electrode, so that the spike waveform changes over the course of the data set; no such drift is seen for the waveforms analyzed in this work. A second mechanism is a change in the shape of the underlying action potential that depends on the history of firing by the cell. In support of this conjecture, we observe that the shape of a spike waveform depends on the interval from the preceding spike, denoted the interspike interval (ISI or τ). For ISIs > 100 ms, the waveform is different, typically narrower, than that for ISIs < 10 ms; this is shown for two single units in Fig. 5, *a* and *c*. Note that we now consider the composite waveform for the stereotrode, $\mathbf{V}^{(n)} \equiv [\mathbf{V}_x^{(n)}, \mathbf{V}_y^{(n)}]$.

To quantify the change in the shape of the waveform, we determined the projection of each spike waveform along a direction defined by the difference between the long-ISI average waveform, denoted $\bar{\mathbf{V}}_L$, and short-ISI average waveform, denoted $\bar{\mathbf{V}}_S$. The projection is defined as

$$P^{(n)} = \frac{(\mathbf{V}^{(n)} - \bar{\mathbf{V}}_L) \cdot \Delta \mathbf{V}_{LS}}{(\Delta \mathbf{V}_{LS})^2} \quad (8)$$

where $\Delta \mathbf{V}_{LS} = \bar{\mathbf{V}}_L - \bar{\mathbf{V}}_S$, so that a value near 0 means a spike waveform has a strong overlap with $\bar{\mathbf{V}}_L$, whereas a projection whose value is near -1 means a waveform has a strong overlap with $\bar{\mathbf{V}}_S$. The n th instance of the waveform has an associated ISI denoted τ_n . A shift in the value of the projection for spike waveforms that follow a long ISI relative to those that follow a short ISI is clearly seen for the two single units of Fig. 5, *b* and *d*, respectively; the shift is comparable with the width of the distribution for either unit. The distribution of projections, integrated over all ISIs, is clearly not Gaussian (Fig. 5*e*).

In general, the time-dependent shift in the shape of the spike waveform is described by a vector, each of whose components is a different function of the ISI. However, the changes we observe can be modeled simply as a constant vector multiplied by a single function of the ISI, denoted $f(\tau)$. The average waveform, denoted $\bar{\mathbf{V}}$, changes according to

$$\bar{\mathbf{V}}(\tau) = \bar{\mathbf{V}}_L + (\bar{\mathbf{V}}_L - \bar{\mathbf{V}}_S)f(\tau) \quad (9)$$

We take $f(\tau)$ to be a single exponential, i.e., $f(\tau) = -\exp[-\tau/\tau_0]$, with $\tau_0 = 22$ ms for the data in Fig. 5*d*. For each instance of the spike waveform, we now calculate the difference between the actual projection (Eq. 8) and the modeled projection, i.e., $P^{(n)} - f(\tau_n)$; the final distribution of these differences is nearly Gaussian (Fig. 5*e*).

Optimal filtering of spike waveforms

We now consider the implications of background variability on the classification of spike waveforms from different

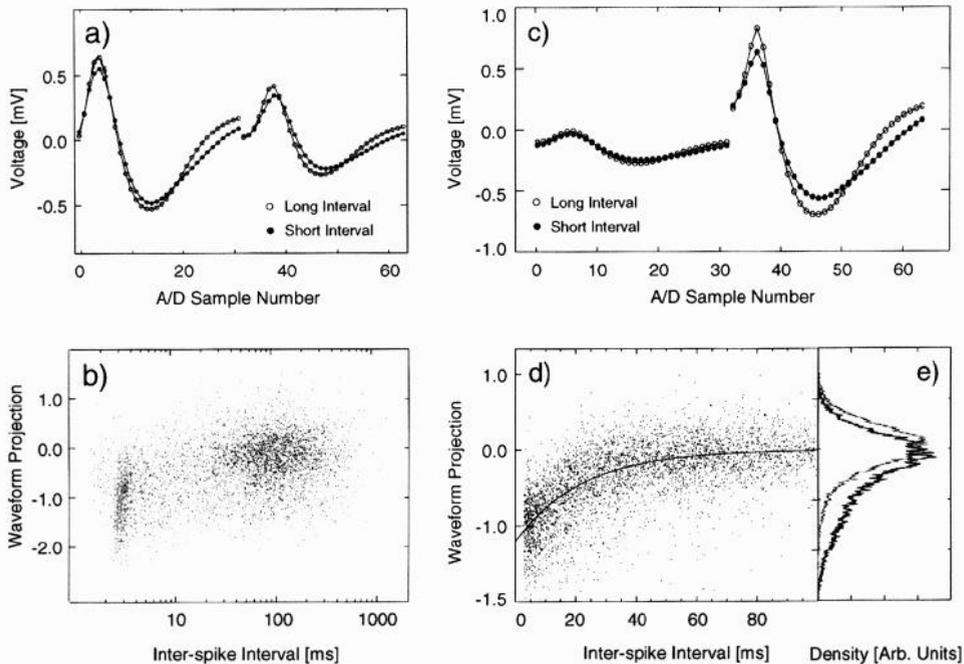


FIG. 5. Systematic variability of spike waveform as a function of the interspike interval (ISI). *a*: stereotyped waveforms observed a short interval after a preceding spike and a long interval after a spike. Note that the short-interval waveform is longer and decreased in amplitude. *b*: distribution of projections of each recorded waveform along the direction $\mathbf{V}_L - \mathbf{V}_S$, defined by the difference between the long- and short-interval waveforms in *a* (see text), as a function of the ISI. Note the evolution in shape as a function of the ISI. *c* and *d*: change in waveform for a 2nd single unit. Solid line through the data: fitted model of the projection to a scalar function of time (Eq. 8). *e*: integrated distribution of projections before and after the subtraction of the model.

single units. Our goal is to construct a filter that accentuates the differences among a set of mean single-unit waveforms (“signal”) in the presence of background activity (“noise”). We first describe the statistical properties of a set of average single-unit waveforms, and then use these properties along with the properties of the background variability to construct this linear filter.

VARIABILITY AMONG SINGLE-UNIT RESIDUALS. We consider a sample of mean single-unit waveforms, denoted $W^{(m)}$, where $m = 1, \dots, M$ labels the mean single unit; in the present case $M = 44$ (22 stereotyped waveform pairs) and each mean waveform is the average of 2,000–5,000 instances that were acquired as segments. Because we are interested in detecting differences among single-unit waveforms, our subsequent analysis is in terms of the difference between a given mean single-unit waveform and the average across all such units (Fig. 6a), denoted $\delta W^{(m)}$, where $\delta W^{(m)} = W^{(m)} - \langle W^{(m)} \rangle$ and the averaging is over the M single units.

The directions of maximum variability among the mean single-unit waveform are given by the principal components of the correlation matrix of the $\delta W^{(m)}$ (Abeles and Goldstein 1977). We denote this matrix C_w , where

$$C_w = \langle \delta W \delta W^T \rangle \quad (10)$$

is of rank T (recall that $T = 32$ and $T < M$) and the average is over the M single units. We observe that the first component is similar to the mean waveform and captures variability in the amplitude of the waveform for different single units (Fig. 6a). Higher-order components do not have simple interpretations, but their largest amplitudes occur in the vicinity of the peak of the spike waveform (Fig. 6a). The ordered spectrum of the corresponding eigenvalues is seen to fall off rapidly with component number (Fig. 6b), such that 95% of the variability

between different single-unit waveforms is accounted for by the four dominant modes shown in Fig. 6a. This result shows that the difference between single units is defined in a subspace whose dimension is substantially lower than that of the original T -dimensional space of the waveforms.

WIENER FILTERING. The space spanned by the principal components of the single units, described above, is a natural basis for the representation of spike waveforms. Following Wiener (Bozic 1994), we seek the optimal linear filter that minimizes the mean square difference between each instance of a spike waveform residual, i.e., a $\delta V^{(n)}$, and the mean single-unit residual that best models that waveform, i.e., one of the $\delta W^{(m)}$ s. We model each instance in terms of the underlying single unit plus noise, i.e.

$$\delta V^{(n)} = \delta W^{(m)} + \mathbf{V}_B^{(n)} \quad (11)$$

where $\mathbf{V}_B^{(n)} = \{V_B^{(n)}(t)\}_{t=1}^T$ is the additive background activity associated with the n th instance of the spike waveform. The filter, denoted \mathbf{F} , is found by minimizing an error, E , defined by

$$E = \langle |\mathbf{F}\delta V - \delta W|^2 \rangle = \langle |\mathbf{F}(\delta W + \mathbf{V}_B) - \delta W|^2 \rangle \quad (12)$$

where \mathbf{F} is a T by T matrix. The average in Eq. 12 is computed over the ensemble of mean waveforms, δW , as well as the ensemble of noise, \mathbf{V}_B . The filter matrix is found by minimizing the error with respect to \mathbf{F} , which gives³

$$\mathbf{F} = C_w(C_w + C_B)^{-1} \quad (13)$$

³The average square error is (Eq. 12) $E = \langle [F(\delta W + V_B) - \delta W][F(\delta W + V_B) - \delta W]^T \rangle = FC_w F^T + FC_B F^T - C_w F^T - FC_w - C_w$, where the correlation matrices C_B and C_w are given by Eq. 10 and 14, respectively, and we note that the averages $\langle V_B \delta W^T \rangle$ and $\langle \delta W V_B^T \rangle$ are 0 under the assumption that the background activity is uncorrelated with the presence of a spike. The filter is found by minimizing E with respect to the filter matrix F , i.e., setting $\partial E / \partial F^T = 0$, from which we get Eq. 13.

where \mathbf{C}_w is given by Eq. 10 and \mathbf{C}_B is the correlation matrix of the background voltage waveforms, i.e.,

$$\mathbf{C}_B = \langle \mathbf{V}_B \mathbf{V}_B^T \rangle \quad (14)$$

where the average is over all instances of background activity (Figs. 2 and 3). The form of \mathbf{F} is considerably simplified when the correlation matrix for the mean single-unit residuals and for the background waveforms diagonalize in the same basis (Eq. 9). In this limit the filter has only diagonal elements in the principal component basis for the mean single-unit residuals,⁴ i.e.

$$(\hat{\mathbf{U}}^T \mathbf{F} \hat{\mathbf{U}})_{\alpha\alpha'} \cong \frac{\lambda_\alpha^2}{\lambda_\alpha^2 + \sigma_\alpha^2} \delta_{\alpha\alpha'} \quad (15)$$

where $\hat{\mathbf{U}}$ is the rotation matrix (Eq. 7) constructed from the principal components of the single-unit residuals, λ_α^2 is the variance for the α th component of these residuals, σ_α^2 is the variance for the background waveforms, and $\delta_{\alpha\alpha'}$ is the Kronecker delta function.

The correlation matrix for the background activity (Figs. 2 and 4) was calculated as above (Eq. 14) and rotated (Eq. 7) into the basis of the mean single-unit residuals. We observe that this matrix, $\hat{\mathbf{U}}^T \mathbf{C}_B \hat{\mathbf{U}}$, is nearly diagonal. The variance terms, σ_α^2 , fall off only slowly with increasing component number (Fig. 6b). Details aside, the essential feature is that the variance of the single-unit residuals decreases much more rapidly than that for the noise, such that signal-to-noise ratio, $\lambda_\alpha^2/\sigma_\alpha^2$, exceeds 1 for only for the first six components. The coefficients for the Wiener filter thus approach a value of 1 for the first few terms and decrease rapidly for the high-order terms (Fig. 6b). In the limit of Eq. 13, the eigenvectors of the filter matrix are equal to those of the correlation matrix \mathbf{C}_s (Fig. 6a); an "exact" calculation of \mathbf{F} yields essentially identical results.⁵ Last, it is instructive to compare instances of the same single-unit waveform before and after filtering, i.e., $\mathbf{V}^{(n)} = \delta \mathbf{V}^{(n)} + \langle \mathbf{V} \rangle$ versus $\mathbf{V}_{\text{filtered}}^{(n)} = \mathbf{F} \delta \mathbf{V}^{(n)} + \langle \mathbf{V} \rangle$ (Fig. 6c). Note that the sharp features of the peak of the spike waveforms are maintained, but that "noise" across all frequency bands is suppressed; this is the essential advantage of filtering in the basis of the principal components.

Segmentation length

The prescription for an optimal filter of the spike waveform is one practical consequence of our analysis of spike waveform variability. A second practical aspect concerns the length of the segmented waveform, which we took to be relatively long in the studies above. A long record will pro-

⁴ The filter in the basis of the principal components of the single-unit residuals $\delta \mathbf{V}^{(m)}$ is found by applying the rotation matrix $\hat{\mathbf{U}}$ to the filter matrix \mathbf{F} (Eq. 13), i.e., $\hat{\mathbf{U}}^T \mathbf{F} \hat{\mathbf{U}} = \hat{\mathbf{U}}^T \mathbf{C}_w (\mathbf{C}_w + \mathbf{C}_B)^{-1} \hat{\mathbf{U}} = \hat{\mathbf{U}}^T \mathbf{C}_w (\hat{\mathbf{U}} \hat{\mathbf{U}}^{-1}) (\mathbf{C}_w + \mathbf{C}_B)^{-1} (\hat{\mathbf{U}} \hat{\mathbf{U}}^{-1})^T \hat{\mathbf{U}} = \hat{\mathbf{U}}^T \mathbf{C}_w \hat{\mathbf{U}} (\hat{\mathbf{U}}^T \mathbf{C}_w \hat{\mathbf{U}} + \hat{\mathbf{U}}^T \mathbf{C}_B \hat{\mathbf{U}})^{-1}$, where we use the fact that $\hat{\mathbf{U}}$ is unitary, i.e., $\hat{\mathbf{U}} \hat{\mathbf{U}}^T = 1$. When \mathbf{C}_B as well as \mathbf{C}_w are diagonalized by the same rotation, so that (Eq. 7) $\hat{\mathbf{U}}^T \mathbf{C}_B \hat{\mathbf{U}} = \Lambda_B^2$ as well as $\hat{\mathbf{U}}^T \mathbf{C}_w \hat{\mathbf{U}} = \Lambda_w^2$, the rotated filter matrix is diagonal, i.e., $\hat{\mathbf{U}}^T \mathbf{F} \hat{\mathbf{U}} = \Lambda_w^2 (\Lambda_w^2 + \Lambda_B^2)^{-1}$, with elements given by Eq. 15.

⁵ The matrix for the background activity \mathbf{C}_B is close to singular and thus the calculation of the \mathbf{F} is ill conditioned. We sidestep this issue by adding a diagonal term $\epsilon \mathbf{I}$ to the denominator in Eq. 14; good convergence is found, with ϵ roughly 0.1 times the mean size of the elements of \mathbf{C}_B .

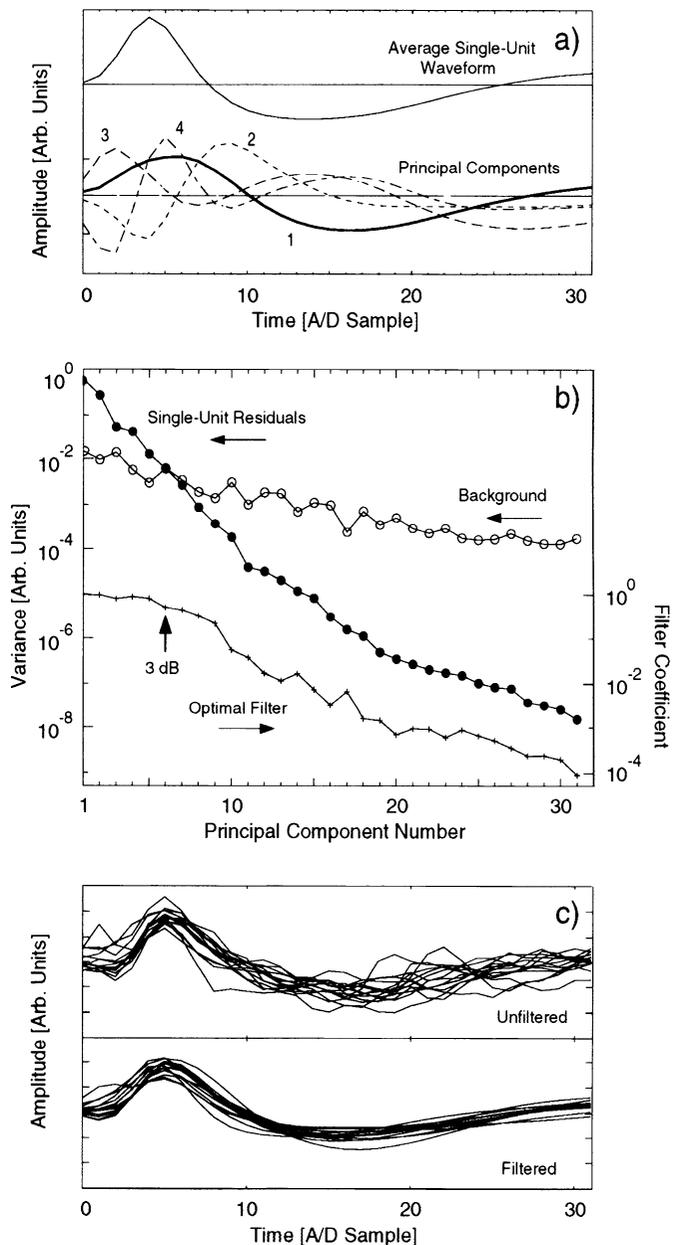


FIG. 6. Variability among a set of single-unit waveforms. *a*: average waveform from a set of 44 single-unit waveforms (top) and the 1st 4 principal components of the residual single-unit waveforms (bottom). *b*: ordered eigenvalue spectrum for the single-unit waveforms and the background activity projected into the basis of the spike waveforms (left scale). The 2 curves cross near mode 6; after this point noise dominates the waveform variability. Bottom trace: optimum (Wiener) filter (Eq. 12). The -3-dB point occurs near mode 6, corresponding to the crossing of the 2 eigenvalue spectra; by mode 9 the amplitude of the filter has fallen an order of magnitude. *c*: Illustration of the optimum filter applied to spike waveforms from the same single unit. In this example the noise level is ~ 4 times the typical level.

vide maximum information about a particular instance of a waveform. On the other hand, too long a record may lead to the presence multiple spikes in the segment, which may confound the sorting process. We use a measure of the information about the mean waveform contained in single-unit waveforms in the presence of noise, i.e., the mutual entropy

between the observed waveform and the underlying ensemble of waveforms, as an objective measure for choosing the desired record length.

The mutual entropy between the observed waveforms, $\delta\mathbf{V}$, and the underlying mean waveform, $\delta\mathbf{W}$, is defined as

$$S(\delta\mathbf{V}, \delta\mathbf{W}) = S(\delta\mathbf{V}) - S(\delta\mathbf{V}|\delta\mathbf{W}) \quad (16)$$

where $\delta\mathbf{V} = \delta\mathbf{W} + \mathbf{V}_B$, as above (Eq. 11), and $\delta\mathbf{W}$ and \mathbf{V}_B will be assumed to have Gaussian distributions with zero mean and covariances \mathbf{C}_W and \mathbf{C}_B , respectively. We use this assumption so as to be able to carry out a calculation based on a realistic amount of data. Note that successive samples must be assumed to be independent for this calculation. As we have illustrated (Fig. 5), this is not completely true; however, to estimate the mutual information we make this assumption. For a multivariate distribution of dimension T with covariance matrix \mathbf{C} , the entropy is given by $S = (T/2) \log_2(2\pi e) + (1/2) \log_2[\det(\mathbf{C})]$ in units of bits (Cover and Thomas 1991). Noting further that $S(\delta\mathbf{V}|\delta\mathbf{W}) = S(\mathbf{V}_B)$ and $\mathbf{C}_V = \mathbf{C}_W + \mathbf{C}_B$ (Eq. 10 and 14), the mutual entropy (Eq. 16) can be written as

$$S(\delta\mathbf{V}, \delta\mathbf{W}) = \frac{1}{2} \log_2[\det(\mathbf{C}_W + \mathbf{C}_B)\mathbf{C}_B^{-1}] \quad (17)$$

In the special case when \mathbf{C}_W and \mathbf{C}_B diagonalize in the same basis, Eq. 18 reduces to $S(\delta\mathbf{V}, \delta\mathbf{W}) \cong (1/2) \sum_{\alpha=1}^T \log_2(1 + \lambda_{\alpha}^2/\sigma_{\alpha}^2)$ and is seen to contain significant contributions only from those dimensions where the variability among the underlying mean waveforms exceeds that of the background, i.e., dimensions for which the ratio $\lambda_{\alpha}^2/\sigma_{\alpha}^2 > 1$.

We calculated the mutual entropy as a function of the length of the segmented record (Eq. 17); the position of the record relative to peak of the spike waveform was adjusted to maximize S . We find that the mutual entropy appears to be close to its asymptotic value for segments with $T = 32$ samples, as used to construct \mathbf{C}_W (Fig. 7). When the segment is decreased to 14 contiguous samples, about a factor of 2 in length, the mutual entropy is reduced by only 1 bit. Increasing the relative amplitude of the background noise, of course, decreases the entropy (Fig. 7).

DISCUSSION

Spike waveforms have at least two sources of variability. First, there are signal sources that persist in the absence of spiking in the observed neuron. To a good approximation, these sources of signal are random with respect to the spiking of the observed neuron and occur at all frequencies, although the power decreases with increasing frequency. Second, there are contributions to waveform variability that are non-random. One such contribution depends on the time since the previous action potential and is likely to result from biophysical changes intrinsic to the observed neuron.

Background variability

The variability of the spike residuals is nearly identical with that of background activity (Fig. 2, *a* and *b*). This suggests that the presence of a spike does not change the average properties of the noise, such as could occur if the

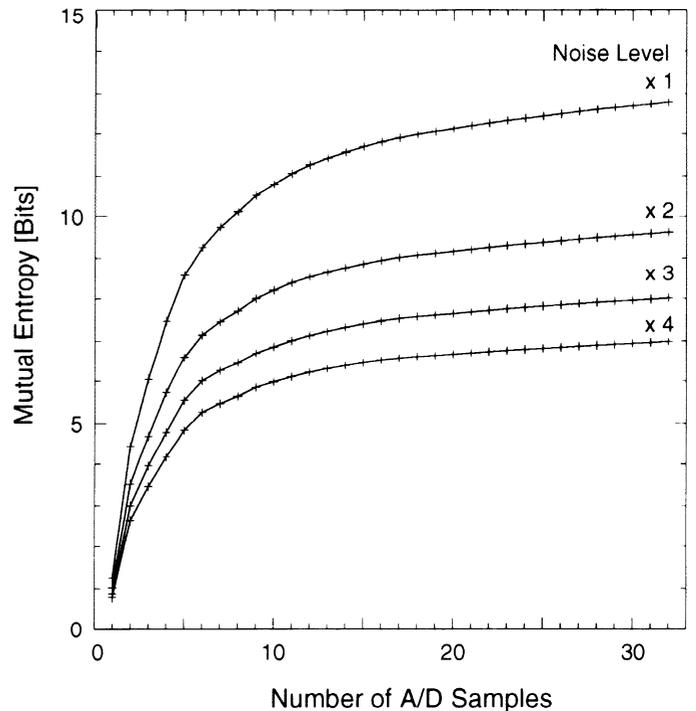


FIG. 7. Entropy of a set of single-unit waveforms, $S(\delta\mathbf{V}, \delta\mathbf{W})$ (Eq. 17) as a function of the length of the segment. Each segment was shifted relative to the peak of the waveform so that the entropy was a maximum. Note that the entropy rises linearly and then rolls over between 5 and 10 samples toward an asymptotic value. Top curve is computed for the data in this work. Bottom 2 curves are computed for noise levels 2 and 4 times those observed.

mean activity in a region was modulated by the spike. Thus the background variability appears as an additive noise.

The power spectrum for the residuals, or background, is not white but rolls off at high frequencies (Fig. 2, *a* and *b*). This shows that there are significant temporal correlations in the spike residuals. Further, the rolloff is slower than that for the spectrum of the spike waveforms. These spectra imply that there are sources of noise other than somatic spikes. Under the assumption that the background consists solely of somatic spikes from an ensemble of neurons, whose arrival times are Poisson distributed on average, the spectrum of the background signal would resemble that of the spike waveform. The observed excess of power at high frequencies in the spectrum of the background activity may result from axons of passage⁶ or fast synaptic currents (Farrant et al. 1994).

We observed that the high-frequency aspect of the variability decayed on a length scale comparable with that between individual wires on the stereotrode pair, $\sim 10 \mu\text{m}$ (Fig. 2, *c* and *d*). This result is consistent with the measured decrement of the amplitude of the extracellular signal with distance for cells in tissue culture, for which the decrement is typically exponential with a decay constant of $\sim 5 \mu\text{m}$

⁶ The power spectrum of propagating action potentials varies as $S(f) \propto f^4$ at high frequencies, whereas that for nonpropagating action potentials varies as $S(f) \propto f^2$. Thus the presence of propagating action potentials in the background activity will boost the high-frequency end of the power spectrum relative to that of mean waveforms, which presumably correspond to spikes at or near somata.

(Tank and Kleinfeld 1986). However, the shape of the extracellular signal and the exact form of the decay with distance depended on the cell type and geometry.

The principal component analysis shows that there is a strong anisotropy in the distribution of spike residuals (Fig. 4). Thus the variability between spike waveforms is much larger along the directions defined by the low principal components than along those defined by the high-order components. Additional anisotropy is introduced by the observed systematic variability in spike waveform as a function of the ISIs (Fig. 5).

A second aspect of the background variability is the presence of 300- to 500-Hz peaks in the spectrum for brief epochs of time, i.e., tens of milliseconds (Fig. 4). The electrical activity associated with these peaks is coherent between wires on the stereotrode pair, i.e., on the 10- μ m scale, but incoherent between different stereotrodes, i.e., on the 1-mm scale (data not shown). One possibility is that interneurons fire trains near their maximal rates for ≥ 10 -ms epochs, for which rhythmic spiking is expected (Gray and McCormick 1996; McCormick et al. 1985). An alternate possibility is that small groups of interneurons fire rhythmically and synchronously for such epochs, although individual neurons in the group may fire at relatively low rates. Last, this aspect of the background variability is suppressed when animals are placed under halothane (2%) anesthesia (unpublished results), not unlike the decrease in the variability of spike arrival times in aroused versus anesthetized or sleeping animals (Paisley and Summerlee 1984).

Spike waveform variability

Our results show that for roughly half of the single units in vibrissa cortex, the individual waveforms evolve as a function of the time since the preceding spike, and reach their asymptotic shape for an ISI of ≥ 100 ms (Fig. 5). This change leads to non-Gaussian distribution of amplitudes in the space of waveforms. Analogous changes in shape are seen in intracellular records for neurons in slice preparations and are particularly strong for cells that produce bursts of spikes, such as layer 5 pyramidal neurons (Connors and Gutnick 1990; McCormick et al. 1985). However, we often see ISI-dependent changes in waveforms that do not exhibit bursting (Fig. 5d).

An important aspect of our analysis is that the change in waveform can be modeled as a linear superposition of two vectors that is parameterized by a single function of time (Fig. 5e). Our analysis suggests that, in principal, changes in the state of a neuron may be inferred from systematic variations in the extracellular signal. It remains to be seen whether such changes in cortical neurons may be related to behaviorally or computationally relevant events.

Implications for spike sorting

ANISOTROPIC VARIABILITY. Our results suggest the importance of correctly accounting for the variability between spike waveforms. In particular, algorithms based on the assumption of an isotropic variability and a Gaussian distribution of amplitudes (Lewicki 1994) are likely to sort a given single-unit cluster into multiple clusters. This problem may

be alleviated by directly modeling the background variability, which is nonstationary (Fig. 3), and the intrinsic waveform variability, such as that associated with the ISI (Fig. 5). A second possibility is to use a hierarchical clustering scheme to account for the anisotropic, non-Gaussian variability (Fee et al. 1996). Application of the latter method to multiunit signals collected from rat primary somatosensory cortex has allowed three or more single units to routinely be classified from a single stereotrode (Fee et al. 1995).

FILTERING. The directions of variability within a cluster that do not lie along the significant directions of variability between different single-unit clusters do not contribute to the discrimination of spike waveforms. Rather these directions only contribute to the total variance of a cluster. Thus the variability in a small number of dimensions contributes to our ability to discriminate between different units. The Wiener filter we describe (Eq. 13; Fig. 6b) preferentially suppresses the variability in directions that are orthogonal to those between different single-unit waveforms.

The application of the Wiener filter (Eq. 9–13) to an instance of a spike waveform, $\mathbf{V}^{(n)}$, follows standard procedures (Bozic 1979). 1) Subtract the mean single-unit waveform, $\langle \mathbf{W} \rangle$, from the waveform to construct the residual, $\delta \mathbf{V}^{(n)}$. 2) Multiply the residual by the filter to form $\mathbf{F} \delta \mathbf{V}^{(n)}$. These vectors may then be clustered, as described in Fee et al. (1996). Recall that relatively few dimensions account for the major fraction of spike waveform variability. The filter may be approximated by keeping only the components whose amplitude is significantly greater than zero, e.g., the first 10 components for the filter in Fig. 6b, so that the filtered residuals may be sorted in a relatively low dimensional space of principal components.^{7,8} We found that this filter typically reduced the total variance of a cluster by a factor of ~ 2 (Fig. 6c).

The eigenvalues of the filter (Fig. 6a) are a basis set for the representation of any spike waveform residual. In this sense, the filtration process we describe builds on the program of Abeles and Goldstein (1977) (see also Gerstein et al. 1983; Gozani and Miller 1994; Roberts and Hartline 1975; Stein et al. 1979) to define an optimum set of functions for the sorting of spike waveforms. Thus different spike waveforms found in different regions of the brain may require different filters. Furthermore, the filter coefficients depend on the signal-to-noise ratio (Eq. 14) in a given recording situation.

SEGMENT LENGTH. We observe that most of the variability between single-unit waveforms occurs near the peak of the spike waveform (Fig. 6a), a segment that is ~ 0.5 ms in duration. The present analysis provides a quantitative measure of gain in discriminability among a set of single-unit spike waveforms that is afforded by the use of longer segments of data. Although longer record lengths certainly improve the discriminability between spike waveforms and do

⁷ An alternate way to view the filter is in terms of the distance metric $\mathbf{F}^T \mathbf{F}$; the metric weights the scalar distance between two waveforms, \mathbf{d}_1 and \mathbf{d}_2 , according to $\mathbf{d}_1^T (\mathbf{F}^T \mathbf{F}) \mathbf{d}_2$.

⁸ A recent application of Wiener filtering to construct a matched filter for individual spike trains (Gozani and Miller 1994) considered filtering in the Fourier frequency domain, rather than in the domain of principal components; in that case there is no reduction in the effective dimensionality of the space of waveform residuals.

not present practical problems as far as storage or computation, they can confound the ability to sort extracellular signals that contain many overlapping waveforms. Our results suggest that records lengths of 1.3 ms afford only a 1-bit improvement in the mutual entropy over a record length half as long (Fig. 7). Thus records containing only the peak region of the waveform may be adequate for sorting spike waveforms, as previously observed (Lewicki 1994).

We thank D. J. Thomson for many useful discussions that aided our analysis, and G. Buzsaki, B. W. Connors, and D. A. McCormick for useful conversations on neuronal firing properties.

Address for reprint requests: D. Kleinfeld, Dept. of Physics 0319, University of California, 9500 Gilman Dr., La Jolla, CA 92093.

Received 4 April 1996; accepted in final form 13 August 1996.

REFERENCES

- ABELES, M. AND GOLDSTEIN, H. M. Multispike train analysis. *Proc. IEEE* 65: 762–772, 1977.
- AHMED, N. AND RAO, K. R. *Orthogonal Transforms for Digital Signal Processing*. Berlin: Springer-Verlag, 1975.
- BOZIC, S. M. *Digital and Kalman Filtering* (2nd ed.). New York: Halsted, 1994.
- CONNORS, B. W. AND GUTNICK, M. J. Intrinsic firing patterns of diverse neocortical neurons. *Trends Neurosci.* 13: 99–104, 1990.
- COVER, M. C. AND THOMAS, J. A. *Elements of Information Theory*. New York: Wiley, 1991.
- FARRANT, M., FELDMEYER, D., TAKAHASHI, T., AND CULL-CANDY, S. G. NMDA-receptor channel diversity in the developing cerebellum. *Nature Lond.* 368: 335–339, 1994.
- FEE, M. S., MITRA, P. P., AND KLEINFELD, D. Central versus peripheral determinates of patterned spike activity in rat vibrissa cortex during whisking. *Soc. Neurosci. Abstr.* 1995.
- FEE, M. S., MITRA, P. P., AND KLEINFELD, D. Automatic sorting of multiple unit neuronal signals in the presence of anisotropic and non-Gaussian variability. *J. Neurosci. Methods*. In press.
- GERSTEIN, G. L., BLOOM, M., ESPINOSA, L., EVANCZUK, X., AND TURNER, M. Design of a laboratory for multi-neuron studies. *IEEE Trans. Sys. Man Cybern.* 13: 668–676, 1983.
- GOLUB, G. H. AND VAN LOAN, C. F. *Matrix Computations* (2nd ed.). Baltimore, MD: Johns Hopkins Univ. Press, 1989.
- GOZANI, S. N. AND MILLER, J. P. Optimal discrimination and classification of neuronal action potential waveforms from multiunit, multichannel recordings using software-based linear filters. *IEEE Trans. Biomed. Eng.* 41: 358–372, 1994.
- GRAY, C. M. AND MCCORMICK, D. A. Chattering cells: superficial pyramidal neurons contributing to the generation of synchronous oscillations in visual cortex. *Science Wash. DC*. In press.
- LEWICKI, M. S. Bayesian modeling and classification of neural signals. *Neural Comput.* 6: 1005–1030, 1994.
- LLINAS, R. R. The intrinsic electrophysiological properties of mammalian neurons: insights into central nervous system function. *Science Wash. DC* 242: 1654–1664, 1988.
- MCCORMICK, D. A., CONNORS, B. W., LIGHTHALL, J. W., AND PRINCE, D. A. Comparative electrophysiology of pyramidal and sparsely spiny stellate neurons of the neocortex. *J. Neurophysiol.* 54: 782–806, 1985.
- MCCAUGHTON, B. L., O'KEEFE, J., AND BARNES, C. A. The stereotrode: a new technique for simultaneous isolation of several single-units in the central nervous system from multiple unit records. *J. Neurosci. Methods* 8: 391–397, 1983.
- MITRA, P. P., FEE, M. S., AND KLEINFELD, D. The variability of extracellular spike waveforms is not random: a new algorithm for spike sorting. *Soc. Neurosci. Abstr.* 1995.
- NATIONAL INSTITUTES OF HEALTH. *Guide for the Care and Use of Laboratory Animals*. Bethesda, MD: NIH Publication 85-23, 1985.
- PAISLEY, A. C. AND SUMMERLEE, A. J. S. Relationships between behavioral states and activity of the cerebral cortex. *Prog. Neurobiol.* 22: 155–184, 1984.
- PERCIVAL, D. B. AND WALDEN, A. T. *Spectral Analysis for Physical Applications: Multitaper and Conventional Univariate Techniques*. Cambridge, UK: Cambridge Univ. Press, 1993.
- ROBERTS, W. M. AND HARTLINE, D. K. Separation of multi-unit nerve impulse trains by a multi-channel linear filter algorithm. *Brain Res.* 94: 141–149, 1975.
- SIMONS, D. Response properties of vibrissal units in rat SI somatosensory cortex. *J. Neurophysiol.* 41: 798–820, 1978.
- STEIN, R. B., ANDREASSEN, S., AND OGUZTORELI, M. N. Mathematical analysis of optimal multichannel filtering for nerve signals. *Biol. Cybern.* 32: 97–103, 1979.
- TANK, D. W. AND KLEINFELD, D. The relation between electrode placement and the amplitude of extracellularly recorded action potentials (Abstract). *Biophys. J.* 49: 223, 1986.
- THOMSON, D. J. Spectral estimation and harmonic analysis. *Proc. IEEE* 70: 1055–1096, 1982.