Chapter 1

Spike Sorting

1.1 Introduction

The point process component of an extracellular recording results from the spiking activity of neurons in a background of physical and biological noise (Section 8.5.1). When a recording electrode measures action potentials from multiple cells, these contributions must be disentangled from the background noise and from each other before the activity of individual neurons can be analyzed. This procedure of estimating one or more single cell point processes from a noisy time series is known as *spike sorting*. When it succeeds, it can transform a fundamental weakness of extracellular recording, namely the inability to isolate changes in the firing rate of single neurons, into one of its greatest strengths - simultaneous measurement from multiple cells (Brown et al., 2004; Buzsaki, 2004).

A range of different approaches is used to address this ubiquitous problem, from manual pattern recognition to more automated, algorithmic solutions (e.g., Abeles and Goldstein, 1977; Lewicki, 1998; Fee et al., 1996a; Nguyen et al., 2003; Pouzat et al., 2004; Quiroga et al., 2004). The algorithmic approaches vary in their assumptions about noise statistics, incorporation of domain knowledge specific to the recording area, and the criteria for identifying single cells. Nevertheless, most of the algorithms can be viewed as different implementations of a common series of steps. This chapter develops a framework for these steps and proceeds to discuss the practical considerations of each level without reference to a specific computational approach. Throughout, the transformations of the data are illustrated by an idealized example modeled on recordings taken from the mammalian retina (Sekirnjak et al., 2006).

1.2   General Framework

In the general case, spike sorting takes a voltage trace containing action potentials from multiple cells and attempts to produce one or more collections of spike times, each corresponding to a putative single cell present in the raw trace. This transformation is accomplished through the following sequence:

1. Acquire raw signal containing extracellular action potentials.

2. Detect candidate neuronal waveforms.

3. Align waveforms to reduce temporal jitter.

4. Calculate shape parameters to simplify spike waveforms, if needed.

5. Identify waveforms with similar shape and segregate them into groups.

6. Judge the quality of these groups.

This transformation from raw signal to collections of similar waveforms can be given a quantitative interpretation as follows. If we consider the signal values in the neighborhood of a data sample, we can define a voltage waveform centered on each time point. We may choose to work directly with the voltage values of these waveforms or we may calculate a reduced set of shape parameters from them to ease visualization or computation. In either case, the waveforms are each associated with a numerical vector and they can be treated as a cloud of points in a high-dimensional space (Section 5.2.1).

The detection of candidate waveforms is equivalent to the segregation of this cloud into points generated from background noise and points that appear to come from action potentials. If the variability in the action potential recorded for a single cell is smaller than the differences in spike shape from different cells, the cloud of spike points will form distinct clusters. The central goal of spike sorting is to identify these clusters of like points in the space of waveforms and use validation measures to judge whether each cluster is likely to have come from an individual neuron.

1.2.1   Manual sorting

A common, and labor-intensive, approach is to reduce each waveform vector to two shape parameters, plot these parameters against one another, and manually identify clusters by visual inspection. This approach has the advantage that the experimenter

can readily apply knowledge specific to the cell population under study. The waveforms, receptive fields, or firing patterns of putative single units can be compared to their known or expected characteristics. A researcher uses this knowledge to refine inclusion criteria and develop confidence that resulting analyses are based on the activity of a single cell.

Manual procedures are, however, of limited utility because (1) shape parameters designed for human inspection are inefficient at representing complex waveforms, (2) the labor-intensive process scales poorly to experiments performed with large numbers of electrodes, and (3) a subjective approach makes it difficult to design reproducible and reportable quality metrics. For these reasons, an algorithmic approach is desirable, and in fact, computational solutions with limited human monitoring have been shown to generally outperform manual sorting (Harris et al., 2000).

1.3   Data Acquisition

The first step in any spike sorting algorithm involves the acquisition of extracellular data in a form amenable to the detection of neuronal spikes.

1.3.1   Multiple Electrodes

An extracellular recording from a local group of neurons begins as an analog voltage signal on each of C electrodes, where C is defined as the number of electrodes in close enough proximity that a single neuron's activity could be recorded simultaneously on each (Section 8.4.2). For example, consider an experiment with two tetrodes placed 2 mm apart, where each tetrode is a four wire electrode bundle with an internal spacing of 50 μm. A spike from a single neuron is unlikely to appear simultaneously on the two bundles due to the large separation between them, but may register on multiple wires within a given tetrode. For purposes of spike sorting, there is no advantage to considering all eight wires as a single source of intermixed spikes, and the signals from each tetrode should be sorted separately. In this case, C would be defined as 4.

Multiple channel electrodes, often stereotrodes ($C = 2$) or tetrodes ($C = 4$), are used to gain an advantage in disambiguating waveforms (McNaughton et al., 1983; Gray et al., 1995). If two neurons with similarly shaped action potentials are equidistant from a single recording electrode, the resulting waveforms will appear indistinguishable. A larger number of recording sites makes it less likely that two neurons will be equally distant from all electrodes. Empirical studies confirm an

increase in waveform discriminability obtained by using multiple electrodes (Harris et al., 2000).

1.3.2   Sampling

In general, the storage and analysis of action potential waveforms requires that the analog data be filtered, sampled and digitized. The data acquisition process, while not specific to spike sorting, introduces several parameters that should be chosen with the sorting process in mind.

The fundamental parameter for acquisition is the sampling rate, F_s, measured in Hz. F_s must be large enough to ensure that multiple samples are recorded during the most rapid voltage swing of the extracellular spike. A rough estimate for a lower bound on F_s can be obtained as follows. The steepest voltage change in an intracellular sodium action potential occurs during depolarization, which lasts approximately 200 μs at physiological temperature. Due to capacitive coupling, the extracellular action potential appears as a time derivative of the intracellular spike (Gydikov and Trayanova, 1986), so the signal will rise and fall in this interval. This suggests 5 kHz as a lower bound on the highest frequencies present in a recorded spike. Empirically, measured extracellular waveforms show significant frequency content at values up to about 8 kHz (Fee et al., 1996b), and the sampling theorem (Section 5.3.5) prescribes a minimum F_s of 16 kHz. Values of 20 kHz to 30 kHz are typically used in practice, exceeding the minimum somewhat to obtain smoother waveforms.

Before sampling, the analog signal must be low-pass filtered with a cutoff frequency below the Nyquist frequency of $F_s/2$ to avoid aliasing (Section 5.3.5). The details of filter selection are beyond the scope of this discussion, but the general requirement is for a multi-pole filter with a cutoff set below $F_s/2$ to account for the roll-off characteristics of real filters. The value used does not directly affect the spike sorting process, but should be noted since it determines the correlation length of the sampled data. For example, data recorded with F_s equal to 30 kHz and low-pass filtered at 5 kHZ will exhibit greater correlation in neighboring samples than if the recording used a low-pass cutoff at 10 kHz. The number of samples over which correlation is significant can be seen in the covariance matrix of the spike waveforms and determines the effective number of degrees of freedom in an N sample waveform.

As an aside, filters with a nonlinear phase response should be avoided because of the resulting distortion of the spike waveforms. While this distortion may not interfere with spike sorting, the resulting spike shapes will depend on the specific filter used and thus difficult be to compare to data taken with different equipment. Linear phase filters minimize this problem.

1.3.3   Data Windows

After acquisition, a high-pass filter is generally needed to isolate action potentials so they can be compared without regard to remaining differences in low-frequency baseline. The guiding factor is the expected duration N of a spike. For example, to ensure that a waveform of duration 1.0 ms will not show a baseline offset due to underlying low frequency activity, frequencies below 1 kHz must be removed from the raw recording.

The estimate N of the duration of a spike also determines the length of the window used to define spike waveforms in the C-dimensional time series obtained from sampling. This window must be chosen so that successive appearances of action potentials from the same cell are as similar as possible from instance to instance. This requires that N span the interval over which spikes tend to be significantly different from background noise. However, setting N too large increases the probability that action potentials from two different cells will occur in the same window, increasing variability. A value of 1.0-1.2 ms is a reasonable compromise for mammalian cortical neurons at physiological temperature.

These data windows associate a vector describing a possible action potential with each time sample t_i as follows. The voltage values in a length N window centered at t_i are extracted from each of the C channels, and these values are collected to obtain a sequence of numbers that can be interpreted as vector with $N \times C$ dimensions. Vectors containing only low amplitude background noise appear in a high density cloud centered at the origin, and artifact events without a repeatable shape appear in regions of low density. Spike waveforms should cluster together and show a clear separation from the central noise cloud.

This perspective is illustrated in Figure 1.1. The raw data shows several action potentials superimposed on background noise, with example data windows of $N = 1.6$ ms. The grey line defines a voltage threshold chosen to highlight candidate spikes (but see the next section). If we consider only those data windows that are below this threshold and describe each with two simple features, the result a large central cloud (Figure 1.1B) representing background noise. Of the data windows that cross threshold, we focus on those where the crossing occurs at the center of the window to avoid considering the same waveform at different offsets. These points comprise a smaller, distinct cluster (Figure 1.1B).
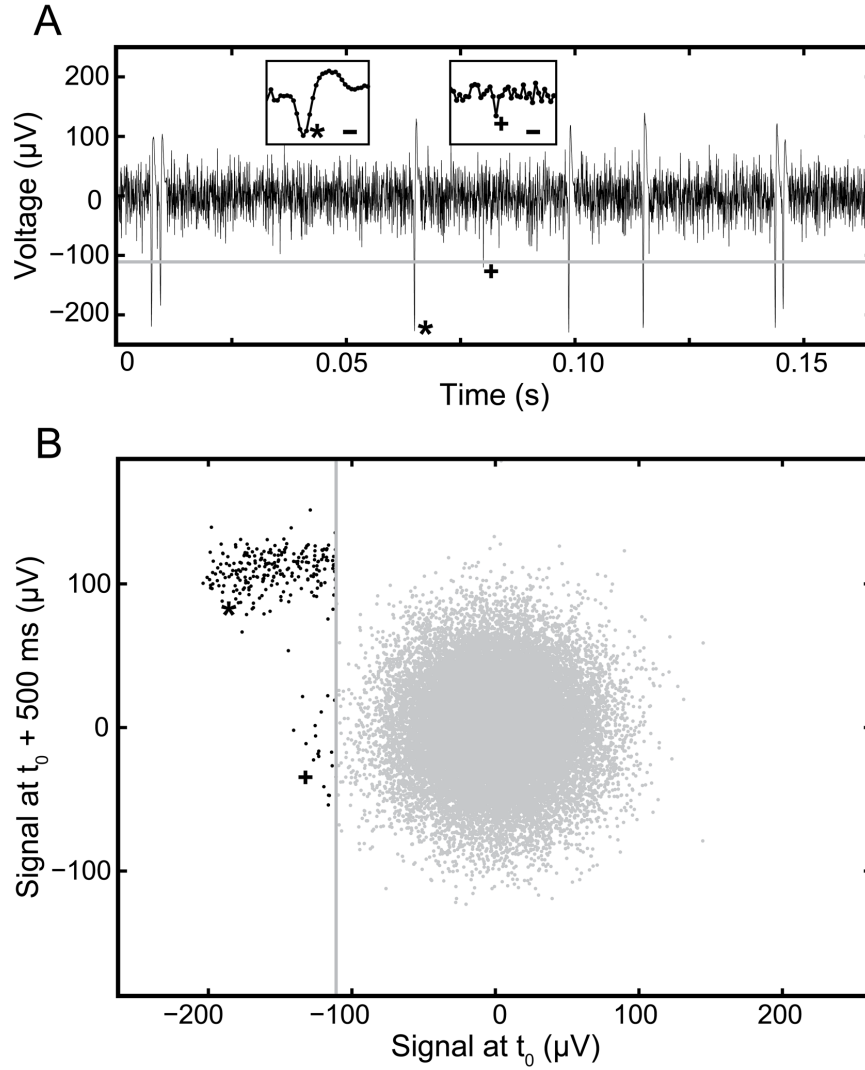
Figure 1.1: Spike and noise events from a simulated data set based on primate retinal ganglion cells. A. Segment of raw voltage time series. Horizontal grey line is the threshold used for spike detection. (*) and (+) indicate a spike and noise event, respectively, that crossed threshold. The insets are expanded views of these events. Inset scale bar is 200 μs. B. Projection of selected windows of length N. Axes indicate voltage values at time relative to center of event, denoted t_0. Vertical line is the threshold for spike detection. (*) and (+) indicate the events marked in A.

1.4 Spike Detection

The extraction of spike waveforms for sorting thus requires separation of data windows associated with neuronal activity from those associated with noise. In this process, a detection function is applied to each data window. The function acts as a possibly nonlinear and time-varying filter takes a length $N \times C$ input vector and reports the evidence for a spike at each time point. Output values above a chosen threshold are taken to be the location of spikes.

When the amplitude of a spike waveform is sufficiently large relative to the thermal and biological background noise, spike detection can be trivial. The detection function might simply reproduce the input signal, and a voltage threshold would be identified to detect large excursions from baseline. However, as the amplitude of spike events approaches the level of noise, a simple threshold includes an increasing number of noise events among the detected spikes. In this case, the statistics of noise and spike waveforms can guide the design of a filter that achieves better separability. For example, Figure 1.1B demonstrates that spike waveforms that are not well separated from noise using a simple threshold may be easily separable in a higher dimension.

A model of background noise can be developed directly from the recorded data (Dimitrov, 2003; Pouzat et al., 2002). Epochs of time known to be free from spiking activity can be used to estimate the parameters of a noise model. For example, if noise is assumed to arise from a Gaussian distribution, the necessary parameters are the mean and covariance. A data window that has a low probability under such a noise model would be considered an action potential. Since biological noise sources are typically non-stationary, the parameters may need to be locally re-estimated throughout the data.

The most direct way to develop a signal model is from a library of prototype spike waveforms. In the presence of white noise, the optimal linear filter for spike detection is a matching filter that convolves the data with the expected waveform shape. If spike sorting is performed in an iterative fashion, the mean waveforms of clusters determined from one round of spike sorting can be used as the matching filters for the next iteration.

When an explicit model of noise or waveform shape is unavailable, a more *ad hoc* approach is necessary. Heuristic approaches are typically based on two observations. First, background noise has zero mean and short auto-correlation. Second, spike events consist of abrupt increases in absolute voltage levels. From these observations, several methods have been developed to detect significant voltage excursions (for a comparison see Obeid and Wolf, 2004). While these techniques may not require the estimation of model parameters, they do not have the optimal properties

of explicit models.

In all of these methods, the choice of threshold for detection of a spike event is a trade off between avoiding false positives (i.e., Type I errors) and false negatives (i.e., Type II errors). A low threshold will capture the most spike events but will erroneously admit many noise events. A high threshold rejects the most noise but also misses the most spikes. When a statistical model of noise is used, the threshold can be calibrated to a desired ratio of Type I to Type II errors. In most cases, a more permissive threshold is desirable because windows of noise can be removed in later stages of processing.

1.4.1 Alignment

Once spike events are detected, the waveform of each action potential is extracted from the data. As discussed previously, a data window is taken from around the time of each spike event. In the case where $C > 1$, this window includes the waveforms from each electrode.

Before the waveforms are clustered, it is important that they are correctly registered with one another. The comparison of voltage values for different instances of an action potential from the same cell is sensitive to the temporal alignment of the two waveforms. It is common to align spike events on the threshold crossing found in spike detection, but this is often inappropriate.

The precise moment at which an analog voltage trace crosses threshold will in general occur between two consecutive samples. The first sample recorded above threshold then lags the actual crossing by a duration between 0 and $1/F_s$ s, and the extracted samples for the entire waveform are equally shifted in time. This process causes two identical waveforms to appear different even in the absence of noise. Background noise compounds this effect by randomly advancing or delaying the actual threshold crossing for each event, possibly by multiple samples.

One solution to misalignment is to digitally upsample the data. When electrophysiological signals are sampled at more than twice the Nyquist frequency, the sampled waveform has all the information needed to reconstruct the continuous analog signal. Interpolation of values between samples increases the effective F_s and yields an alignment with a higher temporal resolution. If the new alignment is based on a single sample, such as a threshold crossing, it will still be sensitive to noise. A more robust alignment point would be a function of the entire waveform, such as its center of mass. After an improved alignment is found, the data can be downsampled to retain the original dimensionality of the spike waveforms. Figures 2A and B show a sample data set before and after alignment, demonstrating the reduction in temporal variability.
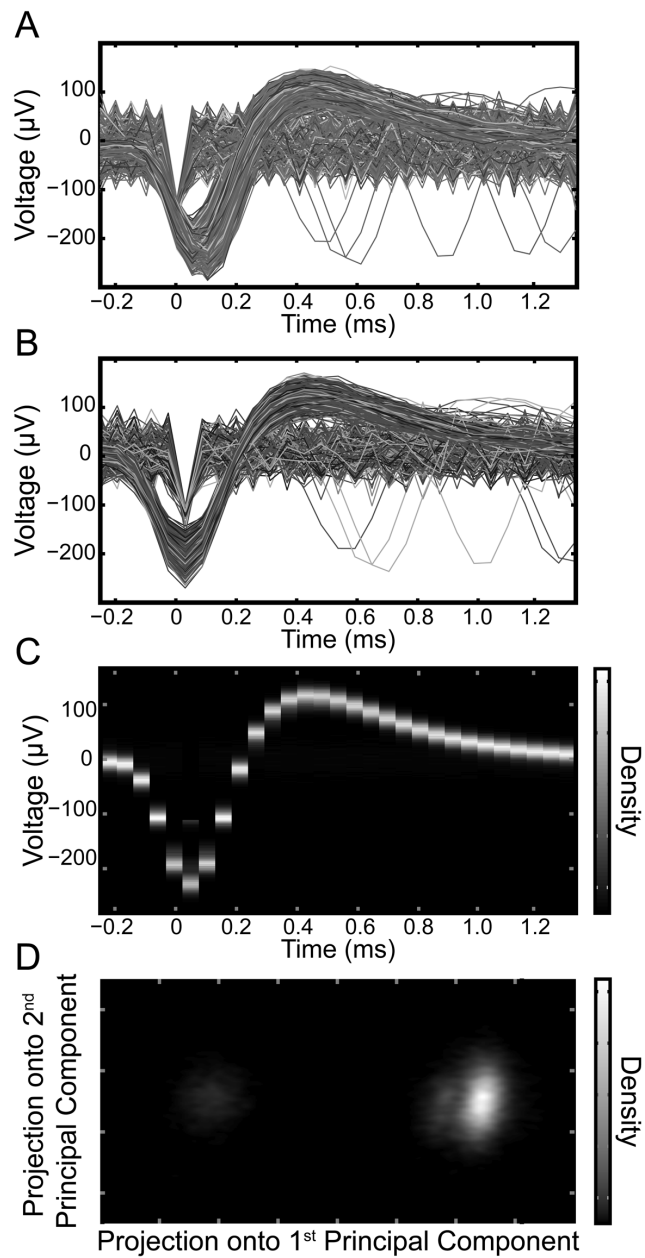
Figure 1.2: Visualization of candidate waveforms. A. Waveforms of all events, before alignment. B. Waveforms of all events, aligned on their peaks. D. Voltage vs. time histogram corresponding to B. D. 2D histogram of waveforms from B projected onto the first 2 principal components of the data set.

1.4.2 Outlier Removal

The detection step produces a set of waveforms that may include outliers, defined here as spike events that are dissimilar to all other waveforms (e.g., Cluster E in Figure 1.3). Outliers may arise from unusual noise events, neurons with very low firing rates, and overlapping waveforms from simultaneously firing neurons. These events are problematic for spike sorting algorithms because they will not appear as dense clusters. Further, many clustering algorithms are not robust to outliers and will distort true clusters in an attempt to account for them. When this is a concern, spikes in low density regions of the waveform space can be set aside as outliers. This can be accomplished, for example, by partially sorting spikes and removing those waveforms that are not well-captured by any cluster.

1.4.3 Data Visualization

Even in automated algorithms, it is useful to visually examine the data to confirm that the results capture the gross structure of the data. This allows the experimenter to understand the degree of separation between signal and noise, the nature of any outliers, and to make corrections to the clustering if necessary.

For example, Figure 1.2B, C and D show three different views of the spikes detected from the data in Figure 1.1 after they have been aligned on their peaks. Apparent in the plot of superimposed waveforms are noise waveforms, such as the one marked by a + in Figure 1.1A, that crossed the detection threshold by a chance fluctuation. Also clear is a group of waveforms that have the expected shape for an action potential, with the variability between waveforms clearly smaller than the separation from noise. A small number of gross outliers can be seen with large secondary peaks.

The superimposed waveforms portray the range of data but can misrepresent the frequency of different events. Figure 1.2C emphasizes common events in a voltage vs. time histogram, where each column corresponds to a histograms of voltage values at a given time point. The background noise vanishes here except for a small peak visible at $t = 0$ because of the noise waveforms show little repeatability. Similarly, the outliers are not visible because of their low relative frequency. The putative spikes, by contrast, have a highly stereotypical waveform. The maximum variability in these waveforms, judged from the voltage spread at each time point, occurs in the samples near the peak of the spike at $t = 0$.

A final perspective on the data comes from treating the waveforms as high-dimensional points, as discussed previously, and projecting them into a two-dimensional subspace (Section 5.2.4). A convenient, though potentially suboptimal, choice of
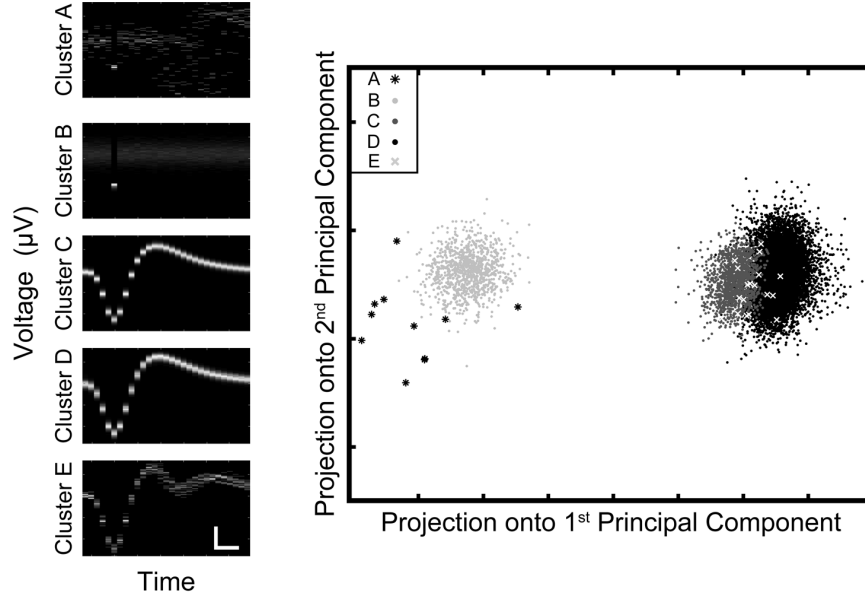
Figure 1.3: Results of spike clustering. Left. Voltage vs. time histogram of the wave-forms of each identified cluster. Scale bar is 100 μV by 200 μs. Right. Waveforms of each cluster projected onto the first 2 principal components.

subspace comes from a principal components analysis (Chapter 7) on the set of spike waveforms. A projection of the data points onto the first two principal components preserves as much of the variation as is possible in an orthogonal two-dimensional representation. When the separation between clusters is significantly larger than the noise within clusters, this maximum variance subspace also approximates the subspace that best preserves cluster separation. Such a projection is shown in Figure 1.2D, plotted as a two-dimensional histogram to highlight regions of high density. The cloud of noise points corresponds to the low density region on the left, and the spike waveforms yield the brighter regions to the right. In particular, this projection suggests that the action potentials are divided into two similar but distinct clusters.

These plots thus demonstrate at a glance what the sorting algorithm should find: two very similar stereotypical waveforms, a background noise cluster, and a small number of outliers.

1.5 Clustering

The core of any complete spike-sorting algorithm defines, either explicitly or implicitly, what is to be considered a cluster of points and divides up the data space

in a way that best fits that definition. For example, a noise model with isotropic Gaussian noise defines a cluster as a compact set of points and an associated mean vector such that points within the cluster are closer in Euclidean distance to their mean than to the mean of any other cluster. The solution for a given number of neuronal sources K, finds the K mean vectors that minimize the average distance from any point to its cluster mean.

Since many physical sources of noise follow a Gaussian distribution, models that assume Gaussian noise are often reasonable and have well-defined solutions. The case in which the noise added to each of K prototypical neuronal waveform is assumed to be from an anisotropic Gaussian distribution is of particular interest because this model has a known optimal solution. As in the previous example, the model defines a cluster so that cluster points are closest to their cluster mean, but in this case the measure of closeness is the Mahalanobis distance. This Gaussian mixture model allows noise that varies unequally or anisotropically over a waveform; this could occur if, for example, the amplitude of the spike peak was less variable than the spike trough due to changing afterhyperpolarization currents.

The full definition of a Gaussian mixture model requires a mean and and a covariance matrix for each of K clusters, as well as mixing parameters describing the relative frequency of the different clusters. Such models are usually solved using an iterative Expectation-Maximization (EM) procedure to estimate maximum likelihood model parameters (Duda et al., 2001). In practice, the main difficulty encountered in this approach lies in determining the number of clusters. Information criteria, such as the AIC or BIC, can find a compromise between goodness of fit and model order by effectively comparing solutions obtained with different numbers of clusters. This approach has been evaluated in an experiment in which simultaneous intracellular and extracellular recordings allowed the results to be validated, and it appears to perform well in many situations (Henze et al., 2000; Harris et al., 2000).

If noise is not Gaussian, for example due to electrode drift or variation in action potential shape due to bursting, then more general definitions of a cluster are needed. Several definitions and techniques are available from the larger literature on clustering (e.g., Duda et al., 2001; Steinbach et al., 2004). One particular notion of a cluster that may be applicable in the non-Gaussian case is the idea of a contiguous set of points with no dips in local density larger than some threshold. This idea is appealing because the threshold allows a hierarchy of clusters, where the predicted number of clusters depends on the amount of discontinuity the experimenter is willing to tolerate. Algorithms such as those designed by Fee et al. (1996a) and Quiroga et al. (2004) have been applied to spike sorting to capture this quality of continuity of density. When the shape of an action potential varies relatively smoothly, as in the slow change of amplitude with electrode drift, these approaches can capture structure missed by Gaussian models.
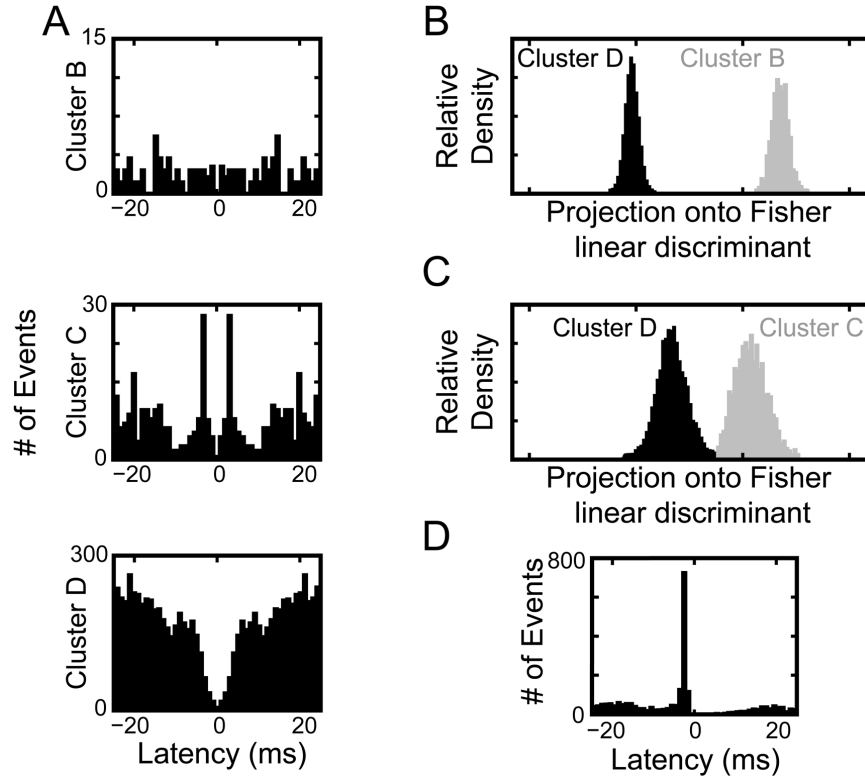
Figure 1.4: Quality metrics applied to spike sorting. A. Spike train autocorrelation for selected clusters. B. Histogram of spike waveforms in cluster B and cluster D when projected onto the Fisher linear discriminant between them. C. Linear discriminant analysis applied to cluster C and cluster D. B. Cross-correlation of spike trains from cluster C and cluster D.

An example applying the algorithm proposed by Fee et al. (1996a) to the data from Figure 1.2 is shown in Figure 1.3. The voltage vs. time histograms for the five resulting clusters are shown on the left and the projection onto the first two principal components is on the right. As predicted, these clusters include a noise cluster, outliers, and two similar neuronal waveforms.

1.6   Quality Metrics

The clusters returned by a sorting algorithm must be evaluated to determine whether they truly represent the output of single cells. Quantitative measures of confidence are necessary to identify clusters suited for single unit analysis. For example, an analysis showing a 5% increase in firing rate should only be performed on units with

an estimated false positive rate of less than 5%.

A common measure of quality is whether the point process associated with a given cluster shows interspike intervals (ISIs; Section 9.5.5) of less than a typical refractory period. Since neurons can not produce action potentials closer together than a few milliseconds, with the details depending on cell type, a single unit cluster will not have ISIs in this range. Conversely, a multi-unit collection of spikes with no biophysical refractory period will have a frequency of short ISIs that is predictable from the mean firing rate of the cluster spikes. The frequency of these refractory period violations, characterized by an ISI histogram or autocorrelation function, will increase as the false positive (Type I) error rate of a cluster goes up. When the mean firing rate is large enough that a significant number of violations are expected by chance, a low frequency of violations is strong evidence that the cluster was generated by a single cell.

While the refractory period provides a measure of multi-unit contamination based on spike times at sufficiently high firing rates, it does not provide an estimate of Type II or false negative errors, corresponding to missed spikes. For example, it is possible to generate clusters without any short ISIs by simply excluding the offending spike. A complementary metric is therefore needed to describe how often similarly shaped spikes are not counted in the same cluster. Although no such metric is currently in common use, the proximity of clusters in the waveform space relative to their internal variability is a natural place to estimate both Type I and Type II errors. An explicit estimate is theoretically possible for parametric models by computing the Bayes' error integral of cluster overlap (Duda et al., 2001). In models that do not form an explicit parametric model for the clusters, nonparametric estimates of cluster overlap may still be valuable for purposes of comparison.

1.6.1 Manual Review

In addition to the quantitative calculation of error estimates, human inspection of the results of a sorting algorithm is needed to confirm the quality of the sorting and to deal with cases in which an algorithm fails. The most common problem arises when an algorithm poorly estimates the number of clusters present in the data, either inappropriately lumping clusters that should be kept distinct or splitting apart collections that seem to be from a single cell.

As an example, Figure 1.4 shows plots used to judge the quality of the clusters from Figure 1.3. Figure 1.4A shows autocorrelation functions indicating that the spikes from clusters C and D both demonstrate refractory periods, since the number of events goes to zero for very short latencies. The noise spikes in cluster B do not appear to show a relative lack of events at these short latencies, although the low

14

frequency of the noise events makes it difficult to accumulate statistics.

In contrast to the variance-preserving principal components analysis, once the data has been divided into clusters it is possible to explicitly find the subspace that best preserves the separation between clusters. In particular, given any two clusters, the Fisher linear discriminant (Duda et al., 2001) defines the one-dimensional projection with the least cluster overlap. Thus Figure 1.4B demonstrates that the action potential cluster B has no overlap with noise cluster D. Figure 1.4C uses the same principle to show that the waveform clusters C and D are in fact separated by a region of low density, despite the high degree of apparent overlap in the principal components projection visible in Figure 1.3. At first glance, it is tempting to conclude that these are then two different cells with similar waveforms.

However, Figure 1.4D shows the cross-correlation of the spike trains for clusters C and D and indicates that spikes from cluster B are very often preceded by spikes from cluster A. This is apparent in the original signal shown in Figure 1.1A, where two short spikes are seen preceded by taller spikes. In this case, knowledge of the recording site is required to understand the cause of this correlation. The retinal ganglion cells that inspired this data set often fire doublet spikes (Sekirnjak et al., 2006), suggesting that cluster B represents a second, smaller spike from the same cell in cluster A. However, it is also possible that the two waveforms represent one cell being driven by another. This evaluation may require human judgment, since the range of factors needed to distinguish these unusual cases is often difficult to define algorithmically. In these cases, a general spike sorting algorithm should compute a best guess solution and present this solution in a way that allows rapid manual confirmation.

Bibliography

Abeles, M. and M.H. Goldstein (1977). Multispike train analysis. *Proc. IEEE* 65(5): 762–772.

Brown, E. N., R. E. Kass, and P. P. Mitra (2004). Multiple neural spike train data analysis: state-of-the-art and future challenges. *Nat Neurosci* 7(5): 456–61.

Buzsaki, G. (2004). Large-scale recording of neuronal ensembles. *Nat Neurosci* 7(5): 446–51.

Dimitrov, Alexander G. (2003). Spike sorting the other way. *Neurocomputing* 52-54: 741–745.

Duda, R.O., P.E. Hart, and D.G. Stork (2001). *Pattern Classification*. Wiley, New York.

Fee, M. S., P. P. Mitra, and D. Kleinfeld (1996a). Automatic sorting of multiple unit neuronal signals in the presence of anisotropic and non-gaussian variability. *J Neurosci Methods* 69(2): 175–88.

Fee, M. S., P. P. Mitra, and D. Kleinfeld (1996b). Variability of extracellular spike waveforms of cortical neurons. *J Neurophysiol* 76(6): 3823–33.

Gray, C. M., P. E. Maldonado, M. Wilson, and B. McNaughton (1995). Tetrodes markedly improve the reliability and yield of multiple single-unit isolation from multi-unit recordings in cat striate cortex. *J Neurosci Methods* 63(1-2): 43–54.

Gydikov, A. A. and N. A. Trayanova (1986). Extracellular potentials of single active muscle fibres: effects of finite fibre length. *Biol Cybern* 53(6): 363–72.

Harris, K. D., D. A. Henze, J. Csicsvari, H. Hirase, and G. Buzsaki (2000). Accuracy of tetrode spike separation as determined by simultaneous intracellular and extracellular measurements. *J Neurophysiol* 84(1): 401–14.

Henze, D. A., Z. Borhegyi, J. Csicsvari, A. Mamiya, K. D. Harris, and G. Buzsaki (2000). Intracellular features predicted by extracellular recordings in the hippocampus in vivo. *J Neurophysiol* 84(1): 390–400.

Lewicki, M. S. (1998). A review of methods for spike sorting: the detection and classification of neural action potentials. *Network* 9(4): R53–78.

McNaughton, B. L., J. O'Keefe, and C. A. Barnes (1983). 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(4): 391–7.

Nguyen, D. P., L. M. Frank, and E. N. Brown (2003). An application of reversible-jump markov chain monte carlo to spike classification of multi-unit extracellular recordings. *Network* 14(1): 61–82.

Obeid, I. and P. D. Wolf (2004). Evaluation of spike-detection algorithms for a brain-machine interface application. *IEEE Trans Biomed Eng* 51(6): 905–11.

Pouzat, C., M. Delescluse, P. Viot, and J. Diebolt (2004). Improved spike-sorting by modeling firing statistics and burst-dependent spike amplitude attenuation: a markov chain monte carlo approach. *J Neurophysiol* 91(6): 2910–28.

Pouzat, C., O. Mazor, and G. Laurent (2002). Using noise signature to optimize spike-sorting and to assess neuronal classification quality. *J Neurosci Methods* 122(1): 43–57.

Quiroga, R. Q., Z. Nadasdy, and Y. Ben-Shaul (2004). Unsupervised spike detection and sorting with wavelets and superparamagnetic clustering. *Neural Comput* 16(8): 1661–87.

Sekirnjak, C., P. Hottowy, A. Sher, W. Dabrowski, A. M. Litke, and E. J. Chichilnisky (2006). Electrical stimulation of mammalian retinal ganglion cells with multielectrode arrays. *J Neurophysiol* 95(6): 3311–27.

Steinbach, M., L. Ertoz, and V. Kumar (2004). The challenges of clustering high dimensional data. In *New Directions in Statistical Physics* (Wille, L.T., editor). Springer-Verlag.