The Pennsylvania State University
The Graduate School

A STUDY OF INTRA-ELECTRODE CORRELATION AND ITS APPLICATION TO NEURAL SPIKE DETECTION

A Thesis in
Electrical Engineering
by
Chinmay Rao

Submitted in Partial Fulfillment
of the Requirements
for the Degree of

Master of Science

August 2011

The thesis of Chinmay Rao was reviewed and approved* by the following:

W. Kenneth Jenkins
Professor of Electrical Engineering
Thesis Advisor, Chair of Committee

Ram M. Narayanan
Professor of Electrical Engineering

Kultegin Aydin
Head of the Department of Electrical Engineering

*Signatures are on file in the Graduate School.

ABSTRACT

Cortical signals recorded during behavioral studies in animals are essential in deciphering the neural basis for adaptation, learning and plasticity. In addition to the procedural challenges involved in implanting micro-electrode arrays, sustaining a chronic recording interface and developing meaningful experimental paradigms there are problems associated with signal acquisition and recording. Animal movement, feeding and grooming actions induce signal artifacts that are temporally and spectrally similar to neural spikes and hence further complicate the task of neural-spike detection and sorting. Here, a simple, first order correlation based technique is presented to clean the recording data of mastication artifact. The method takes advantage of spatial differences between mastication signals and neural signals to identify and eliminate non-neuronal signals from recording data. Signal-to-noise ratios, shape of neural-spike waveform and Receiver Operating Characteristics are plotted to demonstrate the utility of applying correlation algorithm prior to spike detection and sorting. De-noising of recorded signals by this technique improves spike detection efficacy of traditional methods like threshold based detection and Principal Component Analysis (PCA).

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

CHAPTER

1

INTRODUCTION

1.1 Objective and Motivation

Microelectrodes used in cortical implants enable high-resolution in-vivo recording of neuronal activity to understand complex neural processes. The knowledge thus obtained can potentially be used in developing algorithms for future brain-machine interfaces and for understanding the neurophysiological basis of learning and plasticity. In regards to neuroprosthetic developments, implanted microelectrodes also provide a channel for controlled stimulation to combat auditory [1] and visual deficiencies [2], manage Parkinson's disease [3], control epileptic episodes and facilitate communication and movement for the locked-in and quadriplegics [4]. The ability to record from individual neurons using electrodes implanted in the cerebral cortex was demonstrated as early as 1958 [5]. Since then, a variety of electrodes have been developed. Inexpensive, hand-fabricated microwires remain

the electrodes of choice for most of the neural engineering labs. To address the issues of reproducibility and precision, micromachining and photolithography are employed in fabrication of 3-D arrays and planar probes made from silicon, ceramic or polymeric substrate [6, 7, 8, 9]. Each of these arrays are implanted in the brain following standard craniotomy, Once implanted, they transduce extracellular spike activity into voltage signals that are amplified and stored for further analysis.

The detection of neural spike activity is a technical challenge that is a prerequisite for studying many types of brain function. Most neurons in the brain communicate by firing action potentials [10]. These brief voltage spikes can be recorded with a microelectrode, which can often pick up the signals of many neurons in a local region. Depending on the goals of the experiment, the neurophysiologist may wish to sort these signals by assigning particular spikes to putative neurons, and do this with some degree of reliability. In many cases, good single-unit activity can be obtained with a single microelectrode and a simple hardware threshold detector. Often, however, just measuring the activity of a single neuron is a challenge due to a high amount of background noise and because neurons in a local area often have action potentials of similar shape and size. Furthermore, simple approaches such as threshold detection can bias the experiment toward neurons that generate large action potentials. In many cases, the experimental yield can be greatly improved by the use of software spike-sorting algorithms. [10] reviews methods that have been developed for this purpose.

In addition to neural spikes, noise from numerous sources is present in the recorded signals. These sources include: Johnson noise from individual electrodes, field potentials from distant active populations of neurons, EMG signals from jaws, head and entire body, and electrical artifacts generated by sudden animal movements. The situation is exacerbated in behavioral studies wherein the subject is expected to move and perform specific tasks. The detection and sorting of true neural spikes in such a noisy environment remains a challenge that employs schemes to varying complexities. Simple detection schemes usually proceed by threshold isolating signal segments that exceed a predetermined value. Isolated segments are sorted and analyzed depending on specific experimental paradigms using PCA, template matching, wavelet transforms etc. While sorting techniques are interesting in themselves, the subject of this work is improving the detection methods and

evaluating it through enhancement in signal to noise ratios and receiver operating curves. Computationally efficient threshold detection scheme falsely identifies non-neuronal signals like mastication (EMG generated by animal chewing) as being of neural origin. The spatial gradient of these signals results in their near simultaneous appearance on multiple electrodes of an array. Some techniques have been reported that exploit the spatial correlation in noise on the recorded electrodes as means of improving signal-to-noise ratios by removing highly correlated segments [11]. Another approach is to record signals referenced to one of the electrodes in the array instead of a distant ground wire. The first approach works on the entire signal length while the second approach may generate false spike-like signal due to asynchronous cancelation.

One use of spike sorting is to aid the study of neural populations. In some cases, it is possible to measure population activity by using multiple electrodes that are spaced far enough apart so that each can function as a single independent electrode. Traditional methods can then be used, albeit somewhat tediously, to measure the spike activity on each channel. Automatic classification can greatly reduce the time taken to measure this activity and at the same time improve accuracy of those measurements. A further advantage of spike sorting is that it is possible to study local populations of neurons that are too close to allow isolation by traditional techniques. If the activity of several neurons can be measured with a single electrode, it is possible with spike sorting to accurately measure the neural activity, even in cases when two or more neurons fire simultaneously. This capability is especially important for experimental investigations of neural codes that use spike timing.

1.2 Contributions

Here a considerably simple approach is reported that uses correlation only amongst signal segments that are threshold detected as probable spikes on electrodes in an array as a means of eliminating the mastication signals. The approach doesn't require a training set or supervision. It tends to reduce the noise floor by removing large waveforms thereby enabling improved threshold based detection. Initial results point to a reduction in false positives, improvement in SNRs as deduced

from the shape of mean spike waveforms on each electrode. Threshold detection followed by correlation based elimination and signal regeneration can de-noise the electrode data and improve the efficacy of previously reported detection and sorting schemes. One of the sorting schemes studied in detail is Principal Component Analysis. This method is studied and applied on data sets where correlation is minimal, and where a large amount of mastication and interface degradation results in correlated data appearing across different electrodes.

A comparison is presented with respect to objective and subjective measures. Objective measures include a comparison of mean spike waveforms, number of spikes detected, positive phase time span, negative phase time span and peak to peak amplitude of the mean spike waveform. A more subjective tests are the Receiver Operating Characteristics, or the ROC curves, which provides a graphical plot of the fraction of true positives against the fraction of false positives. It shows the tradeoff between sensitivity and specificity (any increase in sensitivity will be accompanied by a decrease in specificity). The closer the curve follows the left-hand border and then the top border of the ROC space, the more accurate the test. However, this test is a subjective test, and the results of the test have been determined by an unbiased examiner, well versed with the nuances of neural spike detection. Finally, the objective measures are re-evaluated, when all tests perform at a probability of equal error.

1.3 Organization

This thesis is organized into five chapters including this one. Chapter 2 presents a review of various spike sorting methods studied by the author, and also the difficulties encountered in Neural spike sorting. Chapter 3 details the data collection methods, the multi-array electrode design, and the algorithm used for spike sorting. Chapter 4 catalogues the results obtained on various data sets, the receiver operating characteristic curves, and other metrics used for the evaluation. Chapter 5 presents all the findings of this study and concludes the report. Two appendices have been included.

CHAPTER

2

REVIEW OF SPIKE DETECTION TECHNIQUES

2.1 Basic Problems in Spike Detection

2.1.1 Multiple Spike Types and Overlapping Spikes

Many of the basic problems in spike sorting are illustrated in the extracellular waveform shown in Figure 2.1

One can immediately observe several different kinds of spikes and it is uncertain whether these correspond to different neurons, or one neuron firing in different ways. Also, there is the problem of overlapping spikes, which is seen better in the example in Figure 2.3. It can be seen that the gradual positive transition of the first spike is overshadowed by the sharp negative transition of the second spike occurring immediately after it. Such occurrences are rare, but need to be accounted for while designing a spike detection technique. Another rare occurrence

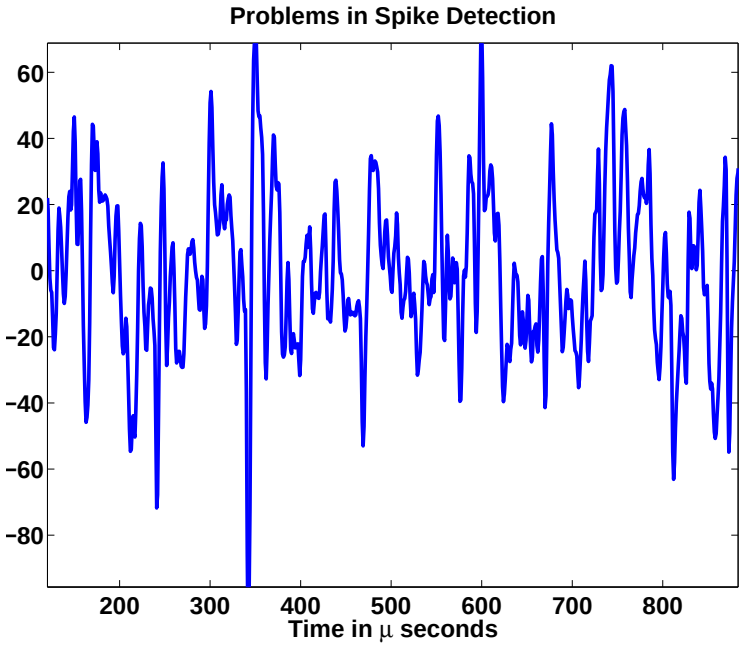

Figure 2.1. An example of a neural waveform

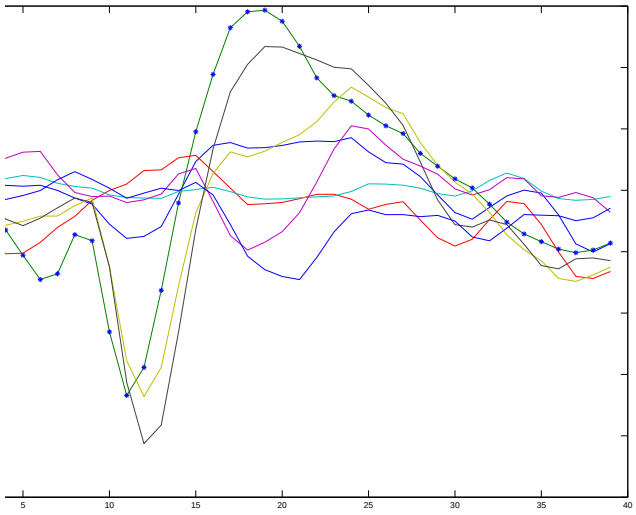

Figure 2.2. An example of a strong neural spike occuring on multiple electrodes

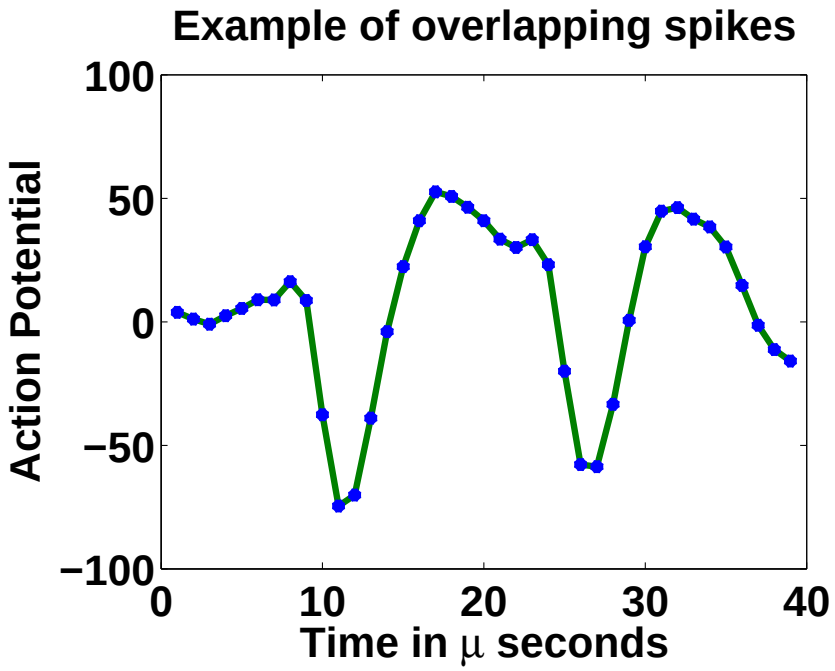

Figure 2.3. An example of overlapping spikes

that needs to be accounted for is shown in Figure 2.2 where a very strong neural signal occurs on several adjacent electrodes.

2.1.2 Electrode Drift

Often during recording the electrode drifts slowly to a new position as the neural tissue settles in response to pressure from the advancement of the electrode. This results in a gradual change in the shapes of the action potentials. This problem can be addressed by the same methods that were used for bursting. Ideally, spike sorters would construct features or templates based on a limited period in the recording and allow these to change slowly over time, so as to compensate for electrode drift. Many of the clustering procedures have the advantage that once a set of classes is determined, the classes can be updated using new data with very little computational expense.

2.1.3 Non-stationary Background Noise

The next problem that is immediately apparent from Figure 2.1 is that of background noise. Noise is present from numerous sources in the recorded signal. These sources include: Johnson noise from individual electrodes, field potentials from distant active populations of neurons, EMG signals from jaws, head and entire body, and electrical artifacts generated by sudden animal movements. The situation is exacerbated in behavioral studies wherein the subject is expected to move and perform specific tasks. Another assumption that can be violated during the course of an experiment is the level of background noise. If the level of background noise remains constant, the classifications will be consistent throughout the trial. If the background noise level fluctuates, there will be many more misclassifications during high levels of noise. Ideally, the estimate of the reliability of the classification should vary with the background noise level, but this is seldom done because of the added complexity of implementing a time-varying noise model.

2.1.4 Variation in Spike Template

The spike height can vary greatly if there are other neurons in the local region that generate action potentials of significant size. If the peak of the desired unit and the dip of a background unit line up, a spike will be missed. This is illustrated in Figure 2.4. It can be seen that all spikes have a sharp negative going transition, but the time phase and peak amplitude vary significantly. Hence, using digital filters, or frequency domain methods becomes difficult.

2.1.5 Mastication Artifacts

The biggest problem encountered during neural spike detection is that of mastication. Electrodes that are inserted into the rat often capture artifacts caused due to mastication movements of the animal. The crucial aspect is that a waveform caused due to mastication looks almost exactly like a spike and this is apparent in Figure 2.5. Almost all traditional single electrode spike detection techniques will detect a mastication signal as a spike signal.

Another example is shown in Figure 2.6

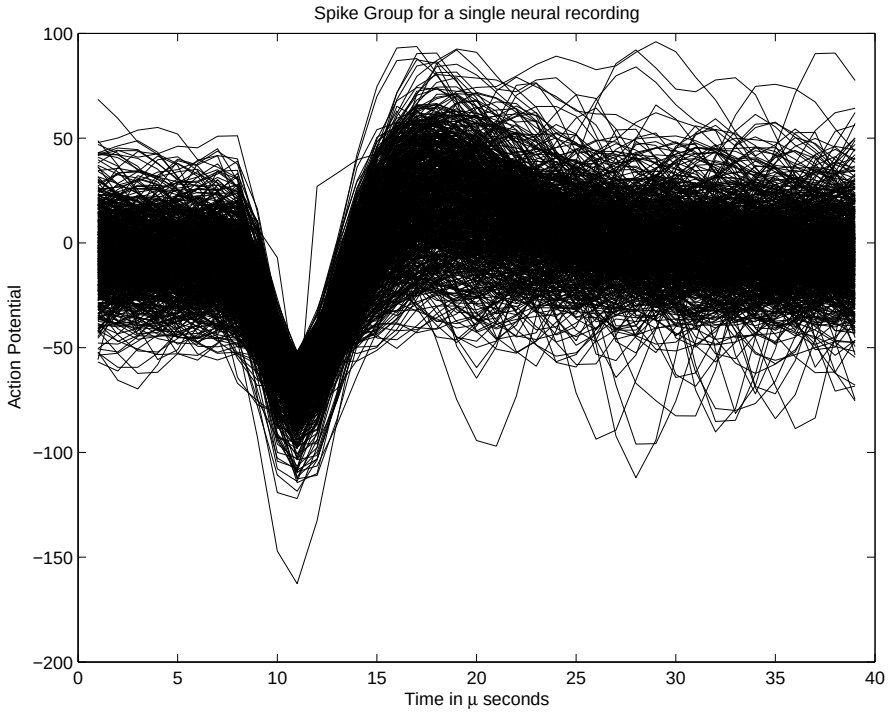

Figure 2.4. Spikes detected during one recording superimposed on each other

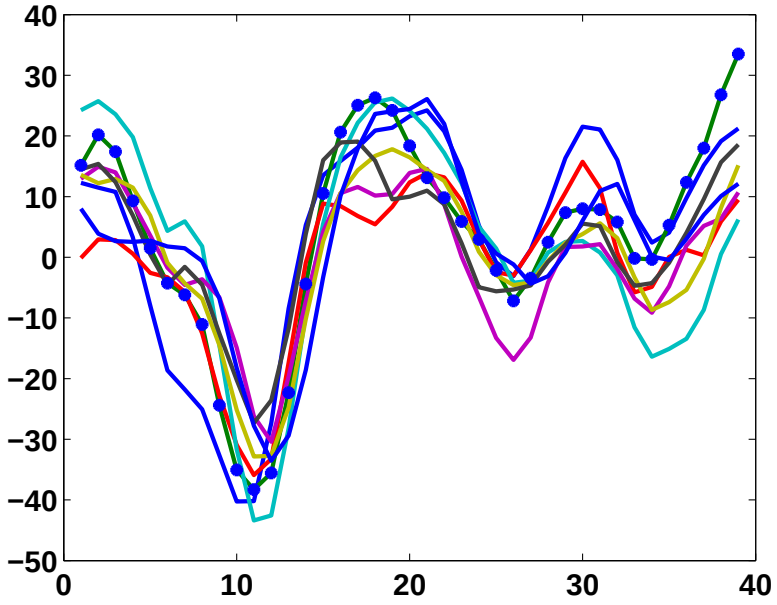

Figure 2.5. A strong mastication signal that resembles a spike, but occurs on all electrodes

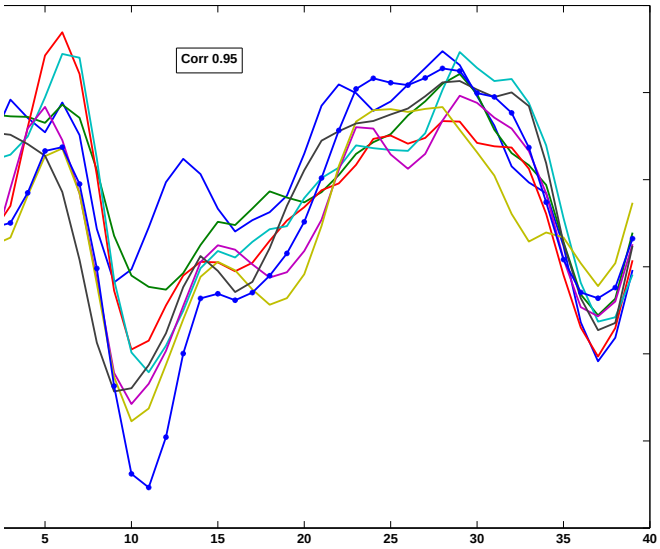

Figure 2.6. A strong spike-like signal that occurs on all electrodes

2.2 Review of Spike Sorting Algorithms

The previous section elaborated some of the difficulties encountered during spike detection. Now, some of the common techniques used to extract spike waveforms are listed.

2.2.1 Feature Extraction

This technique is used both for detecting spikes, as well as classifying them based on their origin. Features such as spike height, width and amplitude are measured and used to differentiate between different waveforms. This is one of the earliest approaches to spike sorting. It was common in these methods to put considerable effort into choosing the minimal set of features that yielded the best discrimination, because computer resources were very limited.

A common method to distinguish between different spike shapes is a technique called cluster cutting. Here, the user defines a boundary for a particular set of features. If a data point falls within the boundary, it is classified as belonging to that cluster; if it falls outside the boundary, it is discarded. The positioning of the boundaries for optimal classification can be quite difficult if the clusters are

not distinct. There is also the same trade-off between false positives and missed spikes as there was for threshold detection, but now in two dimensions. In off-line analysis the cluster boundaries are determined after the data have been collected by looking at all (or a sample from) the data over the collection period. This allows the experiment to verify that the spike shapes were stable for the duration of the collection period. Clustering can also be performed on-line (i.e. while the data are being collected) if the clusters are stable.

A simple approach, such as in nearest-neighbor or k-means clustering, is to define the cluster locations as the mean of data within that cluster. A spike is classified to whichever cluster has the closest mean, using Euclidean distance. This defines a set of implicit decision boundaries that separate the clusters. Classifying in this manner uses only the information about the means and ignores the distribution of data within the cluster. This approach is adequate when the clusters are well separated, but breaks down when clusters significantly overlap or when the cluster shapes differ significantly from a spherical distribution.

2.2.2 Independent component analysis

One recently developed technique that is applicable to multichannel spike sorting is independent component analysis (ICA), which is a method that was originally developed for doing blind source separation[12, 13]. The basic problem is to unmix N independent signals that have been linearly mixed onto N channels with unknown mixing weights. An assumption underlying this technique is that the unknown sources are independent. The objective is to learn the weights that best unmix the channels[14], i.e. transform the mixtures back into independent signals. Under this formulation, the signal separation is performed on a sample by sample basis, i.e. no information about spike shape is used. This is obviously a limitation, but could have advantages in cases where the underlying signals are not stationary.

One restrictive assumption of this approach is that the number of channels (electrodes) must equal the number of sources. This is obviously violated in the case of a single electrode but yields promising results for multichannel optical recordings. Another caution for this method is that it assumes that the sources are mixed linearly. Whether the technique is effective at separating the underlying

signals using multiple electrodes has yet to be tested.

2.2.3 Wavelet Based Methods

Wavelets are scaled and translated copies (known as "daughter wavelets") of a finite-length or fast-decaying oscillating waveform (known as the "mother wavelet"). Wavelet transforms have advantages over traditional Fourier transforms for representing functions that have discontinuities and sharp peaks[15]. In formal terms, this representation is a wavelet series representation of a square integrable function with respect to either a complete, orthonormal set of basis functions, or an overcomplete set of frame functions (also known as a Riesz basis), for the Hilbert space of square integrable functions.

Wavelet transforms are classified into discrete wavelet transforms (DWTs) and continuous wavelet transforms (CWTs). CWTs operate over every possible scale and translation whereas DWTs use a specific subset of scale and translation values.

The typical steps involved in an unsupervised, untrained algorithm using wavelets for spike detection would be as follows:

1. Perform multiscale decomposition of the signal using an appropriate wavelet basis.

2. Separate the signal and noise at each scale.

3. Based on results from the previous two steps, perform Bayesian hypothesis testing at different scales to assess the presence of spikes.

4. Combine the decisions at different scales.

5. Estimate the arrival times of individual spikes.

It is important to choose a wavelet that is suitable for the signal of interest. Common choices are motivated by the shape of the extracellular potentials to be detected in the background noise. This shape can be explained with the help of biophysics. Extracellular and intracellular potentials are intrinsically related. The extracellular potential depends on the transmembrane current, which consists largely of a resistive component and a capacitive component. Since the capacitive

component is proportional to the time derivative of the transmembrane potential, it is approximately biphasic. The capacitive component may dominate the membrane current during an action potential, so the time course of the extracellular spike is typically biphasic. Because a wavelet coefficient represents the measure of similarity between the signal and the basis function, it is reasonable to look for a wavelet that is spike-like. Accordingly, the neural signal would be represented by a few coefficients. In approximation theory this is known as a sparse representation and the basis functions corresponding to these coefficients can be interpreted as intrinsic signal structures. In the presence of noise, the sparse representation becomes an important condition for successful separation of signal and noise. Commonly used wavelets are the gaussian family, db2 and Haar.

Similarly, biophysical considerations of the duration of extracellular potentials can be used to restrict the relevant scales of the wavelet basis functions. Despite their variability in shape and amplitude, the vast majority of extracellular spikes are highly localized in time, with a characteristic duration. For example, action potentials in primate cortex typically last for 0.40.5 ms for signals recorded near the axon, and 0.71.0 ms for recordings near the soma-dendritic complex. Based on this biophysical knowledge of the duration of the transient to be detected, the set of scales for the analyzing wavelet functions can be limited. This practically serves to filter out a significant amount of noise and also appreciably reduce the real-time computational requirements. In summary, a limited set of scales is used.

2.2.4 Interpolation based methods

The clustering methods based on the raw waveform make no assumptions about the spike shape itself. From a Bayesian perspective, this is suboptimal, because it is known a priori that the spike shape will be smooth. What this means in practice is that many more data will be required to obtain a reasonable estimate of the spike shape. If the problem is viewed as one of regularization, there are several approaches that use prior knowledge and obtain more accurate estimates, with much fewer data. This issue is especially relevant in on-line clustering methods where it is important to obtain accurate spike classification based on as few spikes as possible. The objective of regularization (also called noisy interpolation or

smoothing) is to obtain an estimate of an underlying interpolant from noisy data. There is a large literature on this subject.

When combined with the clustering approaches discussed in the previous section, regularization of the spike models places a constraint on the cluster mean so that it represents a smooth function. There is a trade-off between the smoothness of the interpolant and the amount of (additive) noise. If the interpolant is not regularized, it will tend to overfit the data. This is not a very big issue when there is a lot of data, but becomes more important when there are only a few examples. There are many approaches to determining the degree of smoothness. The algorithm of Lewicki [16] uses a simple Bayesian method to determine both the background noise level and the smoothness of the spike functions. This type of model is particularly appropriate for describing action potentials, and it was found that interpolants as accurate as standard regularization models could be obtained using only two thirds the amount of data.

2.2.5 Optimal Filtering Techniques

Another approach to spike sorting uses the methods of optimal filtering. The idea behind this approach is to generate a set of filters that optimally discriminate a set of spikes from each other and from the background noise. This method assumes that both the noise power spectrum and the spike shapes can be estimated accurately. For each spike model, a filter is constructed that responds maximally to the spike shape of interest and minimally to the background noise, which may include the spikes of other units. The neural waveform is then convolved with the set of filters and spikes are classified according to which filter generates the largest response. This is analogous to the clustering methods above in which the metric used to calculate distance to the class mean is inversely related to the filter output. If the filters used to model the spike shapes were orthogonal, this would also be able to handle overlapping action potentials, but in practice this is rarely the case. Comparisons of spike classification on a common data set were carried out by Wheeler and Heetderks who found that optimal filtering methods did not classify as accurately as feature clustering using principal components or template matching.

CHAPTER

3

EXPERIMENTS AND METHOD DESCRIPTION

3.1 Surgical Procedures

All animal procedures followed NIH Guidelines for the Care and Use of Animals and approved by the Penn State IACUC committee (IACUC 20920). Data from two rats that were chronically implanted for 4 weeks for a different study was utilized in this analysis. These subjects were anesthetized with an initial dose of ketamine/xylazine/acepromazine (50:5:1 mg/kg) with additional anesthesia given to maintain areflexia. The subjects were placed in a stereotaxic frame (myNeurolab.com, Product: 463001) and warmed with a heating blanket maintained at 37C. Heart rate and blood oxygen saturation of the animal were monitored with a handheld pulse oximeter (Nonin, 8500AV).

A midline scalp incision was made to expose the cranial plates. Craniotomies were created at 2-4 mm lateral to the midline at sites located 2-4mm anterior to

the bregma. The dura was pierced with a 27G hypodermic needle and further opened with micro-scissors to accommodate the electrode array. A small Gelfoam ®sponge soaked in saline was placed on the sites to maintain hydration. Insertion of electrode arrays into the cortex was done using a three-axis micromanipulator under computer control at a constant speed of 10m/sec. Bone screws were placed in the cranial plates and an acrylic headcap was applied to hold the electrode connectors.

3.2 Neural Recording and Data Analysis

A commercial multi-channel acquisition system (Tucker-Davis Technologies Inc., http://www.tdt.com) was used to collect simultaneous neural recordings from the electrodes of chronically implanted animals. After a unity gain buffering stage, a lightweight bio-amplifier (Medusa, TDT) digitized the signals (25 kHz) with low-noise 16-bit A/D converters. The signals were multiplexed and transmitted via a fiber-optic cable to rack-mounted DSP modules. Signals were digitally band-pass filtered 300-5kHz, down-sampled at 12 kHz and streamed to disk for offline analysis. Recordings were carried out in regular sessions with the animal either awake or lightly sedated with isoflurane. During each session, a 5-minute block of recordings were made from each of the implanted arrays. The electrode impedances were also measured at 1 kHz with a Bak electrode impedance meter.

3.3 Electrode Fabrication and Insertion

The electrodes were composed of eight 50m tungsten microwire insulated with polyimide (California Fine Wire) arranged in a 2x4 array. The recording tips were blunt and separated by 250 m with islands of dental acrylic similar to [24]. The microwires were connected to a small custom PCB interface board which was soldered to an Omnetics nano-miniature connector. A stainless steel ground wire was tied to one of the bone screws to serve as ground (for chronic experiments). Epoxy was applied to reinforce the microwire attachment to the connector. Electrodes used in chronic experiments were sterilized with 5 Megarad dose of gamma radiations.

A schematic diagram of the insertion force measurement system is depicted in Fig. 2. The electrode arrays were attached to the load-cell using a light-weight, custom built flexiglass connector. The load-cell was calibrated to produce an output of 5V for a force of 245.25mN (25gms). Despite working in the lower 5% portion of the dynamic range, the sensor was tested with calibrated weights and found to be linear and accurate down to a load of 0.049mN (5mg). The insertion force experienced by the penetrating electrode arrays attached to the sensitive load cell was measured by coupling the load-cell output to an in-line amplifier (Honeywell Sensotec model MBL 25gms: 0.1%FS linearity; 0.03% FS repeatability and model UV, respectively). The analog output of the in-line amplifier was connected to a digital multimeter for qualitative observation. The output signal of the in-line load cell amplifier was also acquired and stored on a PC using Datawave SciWorks acquisition software for subsequent analysis.

The entire assembly consisting of the load-cell and the microwires was manually lowered toward the craniotomies using a 3-axes micromanipulator. Once a reading was observed on the digital multimeter, the electrodes were retracted by approximately 50m to ensure that array was close but not touching the brain. For insertion, a custom-built computer controlled stepper-motor setup was used to drive an oil hydraulic micromanipulator (Narishige MWO-3). This action directed the microwire array at a constant speed of 10m/sec into the cortex. The final depth of insertion for acute studies was 2mm to ensure complete penetration of the electrodes. The final depth for the chronic studies was restricted to 1mm to enable recording from cortical layer V of the primary motor cortex. The analog output signal of the in-line amplifier was sampled at 50000 samples/second and was stored on a PC using Datawave SciWorks acquisition software for subsequent analysis. An average value was obtained for every 50000 sample points. The final force curves had one data point per second. This smoothened out the effects of any unwanted vibrations caused by air-flows or surgical table movement. The force values were plotted as a function of distance. First maximum of the force-distance curve was considered to be the peak force of insertion. Reduced depth insertions during the chronic studies resulted in the absence of a distinct first-maximum (except R47). Hence maximum force experienced by the electrode array during the insertion run was considered for comparison. Average force during insertion was

calculated from the force-distance curve as well.

3.4 Algorithms Developed for Multi-Electrode Spike Detection

This thesis describes three algorithms used for Multi electrode spike detection. The first is a single electrode thresholding algorithm which has provisions to overcome the overlapping spike problem described in Chapter 2.The second algorithm is the crux of this thesis and is a multi-electrode correlation algorithm, that works on the output of the previous thresholding algorithm. The third algorithm is a Principal Component Analysis algorithm, which again works on one electrode at a time. These methods were chosen since they require the least amount of training to be implemented by someone not familiar with the nuances of their methodologies. Each algorithm is detailed below.

3.4.1 Threshold detection

For many neurons, the most prominent feature of the spike shape is its amplitude, or the height of the spike[17]. One of the simplest ways to measure the activity of a neuron is with a voltage threshold trigger. The recording electrode is positioned so that the spikes from the neuron of interest are maximally separated from the background activity. Neural activity is then measured with a hardware threshold trigger, which generates a pulse whenever the measured voltage crosses the threshold. This method is by far the most common for measuring neural activity. The obvious advantages of threshold detection are that it requires minimal hardware and software, and can often obtain exactly the information the experimenter wants. The disadvantage is that it is not always possible to achieve acceptable isolation.

The simple threshold detection scheme has been enhanced by the addition of several features to overcome some of its limitations:

1. Severe mastication artifacts have a very large negative going transition, or a very high positive phase, that is greater in magnitude than most neural signals are capable of reaching. Thus, at the first stage, for each electrode of

an array, regions of large signal amplitude ($\geq$ 300V or $\leq$ -300V) are threshold detected and eliminated from the signal.

2. For each excursion below the 2.5σ threshold, the minimum point is picked from this excursion. This represents the solution to the overlapping spike problem. Thus, two spikes are considered separated if the first occurring spike transitions away from the 2.5σ region of the signal.

3. A fixed length window is then constructed around each peak detected above. This accounts for the pre-spike fall time, and the subsequent settling time of the spike.

4. A spike bin is created consisting of all spikes, and the corresponding signal from the other electrodes and is stored for further analysis.

If consecutive samples showed 2.5σ excursions, only the minimum of these samples was considered and used for identification of 3ms signal segment. Another important feature of this algorithm is that severe mastication artifacts are removed before the threshold is determined. This results in an improved detection scheme. A block diagram of this stage is shown in Figure 3.1.

3.4.2 Multi-electrode Correlation Algorithm

The next stage of the algorithm is to look at the spike blocks generated in the thresholding algorithm and study the intra-electrode correlation. Few methods have been developed that look at different electrodes to clean the signal [18], or to improve the signal to noise ratio [11]. This section describes the algorithm used to actually detect spikes from the above spike block based on the intra-electrode correlation.

Correlation indicates the strength and direction of a linear relationship between two random variables. In general statistical usage, correlation or co-relation refers to the departure of two variables from independence. A number of different coefficients are used for different situations. The best known is the Pearson product-moment correlation coefficient, which is obtained by dividing the covariance of the two variables by the product of their standard deviations. Despite its name, it was first introduced by Francis Galton.

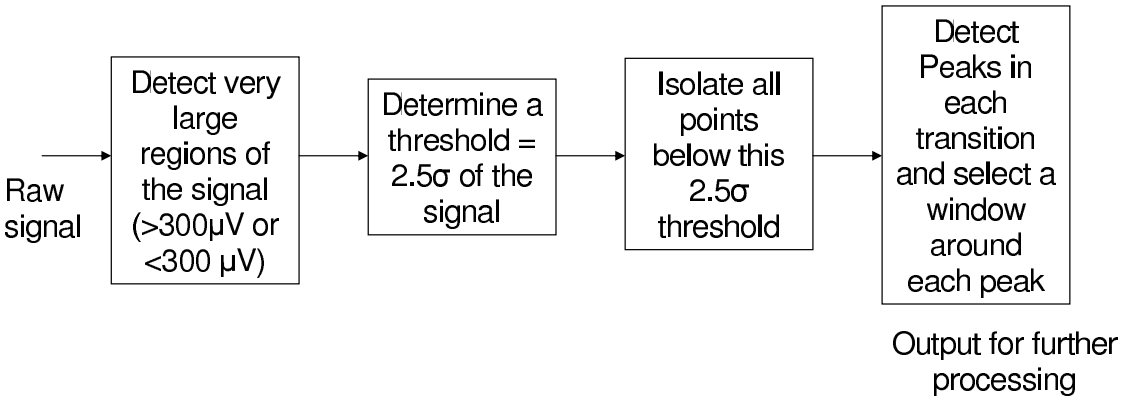

Figure 3.1. A block diagram representing the thresholding algorithm

The correlation ρ_{XY} between two random variables X and Y with expected values μ_X and μ_Y and standard deviations σ_X and σ_Y is defined as:

$$\rho_{X,Y} = \frac{\operatorname{cov}(X,Y)}{\sigma_X \sigma_Y} = \frac{E((X-\mu_X)(Y-\mu_Y))}{\sigma_X \sigma_Y} \tag{3.1}$$

where E is the expected value operator and cov means covariance. Since $\mu_X = E(X)$, $\sigma_X^2 = E(X^2) - E^2(X)$ and likewise for Y, we may also write

$$\rho_{X,Y} = \frac{E(XY) - E(X)E(Y)}{\sqrt{E(X^2) - E^2(X)}\ \sqrt{E(Y^2) - E^2(Y)}} \tag{3.2}$$

The correlation is defined only if both of the standard deviations are finite and both of them are nonzero. It is a corollary of the Cauchy-Schwarz inequality that the correlation cannot exceed 1 in absolute value.

The correlation is 1 in the case of an increasing linear relationship, -1 in the case of a decreasing linear relationship, and some value in between in all other cases, indicating the degree of linear dependence between the variables. The closer the coefficient is to either -1 or 1, the stronger the correlation between the variables.

The covariance matrix is a matrix of covariances between elements of a vector. It is the natural generalization to higher dimensions of the concept of the variance of a scalar-valued random variable.

If entries in the column vector

$$X = \begin{bmatrix} X_1 \\ \vdots \\ X_n \end{bmatrix}$$

are random variables, each with finite variance, then the covariance matrix Σ is the matrix whose Σ_{ij} entry is the covariance

$$\Sigma_{ij} = \mathrm{E}\left[(X_i - \mu_i)(X_j - \mu_j)\right] \tag{3.3}$$

where$\mu_i = \mathrm{E}(X_i)$ is the expected value of the ith entry in the vector X. In other words,

$$\Sigma = \begin{bmatrix} \mathrm{E}[(X_1 - \mu_1)(X_1 - \mu_1)] & \mathrm{E}[(X_1 - \mu_1)(X_2 - \mu_2)] & \cdots & \mathrm{E}[(X_1 - \mu_1)(X_n - \mu_n)] \\ \mathrm{E}[(X_2 - \mu_2)(X_1 - \mu_1)] & \mathrm{E}[(X_2 - \mu_2)(X_2 - \mu_2)] & \cdots & \mathrm{E}[(X_2 - \mu_2)(X_n - \mu_n)] \\ \vdots & \vdots & \ddots & \vdots \\ \mathrm{E}[(X_n - \mu_n)(X_1 - \mu_1)] & \mathrm{E}[(X_n - \mu_n)(X_2 - \mu_2)] & \cdots & \mathrm{E}[(X_n - \mu_n)(X_n - \mu_n)] \end{bmatrix}. \tag{3.4}$$

The correlation of each identified segment on each electrode with respect to simultaneous segments on remaining electrodes of the array was calculated. A correlation matrix was constructed according to equation 3.4 and the diagonal elements were considered. Thus for an eight electrode array, seven correlation coefficients were obtained for each probable spike on each electrode. If any one of the seven coefficients exceeded a predetermined threshold, the spike was considered

to be of non-neuronal origin and hence eliminated from further analysis.

3.5 Principal Component Analysis

One method for choosing features automatically is with principal component analysis [19, 20, 21]. The idea behind principal component analysis (PCA) is to find an ordered set of orthogonal basis vectors that capture the directions in the data of largest variation [22]. The data are the original spikes from the recorded waveform.

Each waveform is centered around the spike maximum, with a fixed window on each end of the spike to minimize the variability of the spike shapes. To represent any particular data point (i.e. a spike) the principal components are scaled and added together. The scale factor for each component is sometimes called the score. PCA involves the computation of the eigenvalue decomposition or Singular value decomposition of a data set, usually after mean centering the data for each attribute.

PCA is mathematically defined as an orthogonal linear transformation that transforms the data to a new coordinate system such that the greatest variance by any projection of the data comes to lie on the first coordinate (called the first principal component), the second greatest variance on the second coordinate, and so on.

PCA can be used for dimensionality reduction in a data set by retaining those characteristics of the data set that contribute most to its variance, by keeping lower-order principal components and ignoring higher-order ones. Such low-order components often contain the most important aspects of the data. But this is not necessarily the case, depending on the application.

For a data matrix with zero empirical mean $\bar{X}$, (the empirical mean of the distribution has been subtracted from the data set), where each row represents a different repetition of the experiment, and each column gives the results from a particular probe, the PCA transformation is given by:

$$\mathbf{Y^T} = X^T\mathbf{W} = V\mathbf{\Sigma} \tag{3.5}$$

where $V\Sigma W^T$ is the singular value decomposition of X^T.

Now, assume that each column is made up of results for a different spike from the bins obtained after correlation, and each row the results from a different sample of the spike. This will mean that the PCA for the data spike matrix will be given by:

$$Y = \mathbf{W}^T X = \mathbf{\Sigma} V^T \tag{3.6}$$

The goal is to choose those features that allow pattern vectors belonging to different categories to occupy compact and disjoint regions in the m-dimensional feature space. The effectiveness of the representation space, generated from the feature set, is determined by how well patterns belonging to different classes can be separated. Given a set of training patterns from each class, the objective is to establish decision boundaries in the feature space to separate patterns belonging to different classes. In the statistical decision-theoretic approach, the decision boundaries are determined by the specified or learnt probability distributions of the patterns belonging to each class.

Feature extraction methods determine an appropriate subspace of dimension d in the original feature space of dimension m with $d \leq m$. The best known linear feature extraction technique, Principal Component Analysis (PCA) that makes use of Karhunen-Love transformation [22] to compute the d largest eigenvectors of the covariance matrix of the N patterns, each of which is m-dimensional. Since PCA uses the most expressive features (eigenvectors with the largest eigenvalues), it effectively approximates the data on a linear subspace using the mean squared error criterion.

CHAPTER

4

RESULTS

4.1 Algorithm Implementation

The subsequent analysis is carried out using Matlab (Mathworks, MA). For each electrode of an array, regions of large signal amplitude ($\geq 300\mu$V or $\leq -300\mu$V) are threshold detected and eliminated from the signal. Points that exceed negative 2.5σ of the signal are detected. A threshold of 2.5σ was selected instead of 3σ to increase the probability of detecting neuronal spikes in the signal, especially in light of inflated signal energy due to the presence of mastication artifacts. 1ms signal prior and 2ms signal after the occurrence of 2.5σ excursions were identified for further analysis. If consecutive samples showed 2.5σ excursions, only the minimum of these samples was considered and used for identification of 3ms signal segment.

Thus in the first stage, probable spike bins were generated for each electrode of the array. The correlation of each identified segment on each electrode with respect to simultaneous segments on remaining electrodes of the array was calculated. Thus for an eight electrode array, seven correlation coefficients were obtained

for each probable spike on each electrode. If any one of the seven coefficients exceeded a predetermined threshold, the spike was considered to be of non-neuronal origin and hence eliminated from further analysis. A careful look test was done to calculated probability density functions. This particular test was subjective in nature wherein a signal segment was classified as spike-like or not and its maximum correlation coefficient was noted down. Probability distribution plots were generated with correlation coefficients on the x-axis and number of signal segments on y-axis. A receiver operating characteristic curve was also generated from this data. Considering correlation coefficient as a discrimination threshold a plot of probable true spikes versus probable false spikes produced the desired ROC.

Correlation coefficients were also calculated by taking the electrode geometry into consideration. For a given electrode, correlation coefficients with other electrodes were multiplied by weighting factors based on their distance from the electrode under consideration. This approach reduced the possibility of eliminating those waveforms wherein a strong spike was registered on multiple closely spaced electrodes.

Subsequent analysis was handled in Matlab (Mathworks, MA). Signals were pre-processed to eliminate large mastication artifacts (¿300 μ V). Probable neural spikes were threshold-isolated and extracted as waveform segments (3 ms in duration) exceeding 3 standard deviations of the raw signal. Further elimination of non-neuronal waveforms took advantage of the spatial difference between neuronal and non-neuronal spikes. It was assumed that non-neuronal spikes occur simultaneously on all the electrodes thereby exceeding a predetermined correlation threshold. In contrast, neuronal spikes are restricted to a couple of closely spaced electrodes in an array. Thus correlated spikes above a set threshold were considered non-neuronal and hence eliminated from the recorded signal that was used for subsequent analysis. Recording performance was evaluated using multiple matrices. A mean-spike was calculated from all the detected spikes on the electrode. Noise signal was computed by eliminating all the detected spikes from the recorded signal. SNR was computed by considering the peak-to-peak value of the mean-spike as signal and twice the 3-standard deviation of the noise signal as noise. Presence of healthy neurons in the proximity of the recording sites can estimate the interface quality. Number of recorded spikes was determined to provide

a rough measure of the interface quality.

PCA was also implemented in MATLAB and works both on the spike blocks detected by thresholding, as well as the blocks detected by the correlation algorithm. Also, a re-thresholding is performed after correlation detects the spike segments that show a high degree of mastication.

The following methods have been implemented and compared

1. Threshold based detection
2. Thresholding followed by correlation
3. Thresholding, followed by correlation, followed by re-thresholding.
4. Principal Components applied to the output of Thresholding
5. Principal Components applied to the output of Thresholding followed by correlation
6. Principal Components applied to the output of Thresholding, followed by correlation, followed by re-thresholding

4.2 Evaluation Measures

Several methods were used to evaluate each algorithm described above. These range from number of spikes detected, various kinds of signal to noise ratios, to probability of detection. The methods have been divided into objective and subjective measures. Since the whole process of qualifying what constitutes a neural spike and what does not is a grey area, measures like probability of detection and receiver operating characteristics come under subjective measures.

4.2.1 Objective Measures

The following objective measures have been considered, which do not require a human subject to evaluate, and can be calculated analytically on a computer:

1. **Number of Spikes Detected:** This represents the number of spikes detected in a fixed amount of data.

2. **Peak Negative Time Phase:** This represents the time elapsed in the negative transition of the mean spike, between the 10% points at the upward and downward transition. The 10% is measured with respect to the peak negative value.

3. **Peak Positive Time Phase:** This represents the time elapsed in the positive settling of the mean spike, between the 20% points at the upward and downward transition. The 20% is measured with respect to the peak positive value.

4. **Peak to Peak Amplitude:** This is the difference between the peak positive and the peak negative amplitude of the mean spike in micro-volts.

5. **Signal to Noise ratio:** Several signal to noise ratios are described in Appendix. These are measured before and after the spike detection, and used as a standard method for evaluation [10]

4.2.2 Subjective Measures and Receiver Operating Characteristics

The Receiver Operating Characteristics, or the ROC curves, which provides a graphical plot of the fraction of true positives against the fraction of false positives. It shows the tradeoff between sensitivity and specificity (any increase in sensitivity will be accompanied by a decrease in specificity). There are four possible outcomes from a binary classifier. If the outcome from a prediction is p and the actual value is also p, then it is called a true positive (TP); however if the actual value is n then it is said a false positive (FP). Conversely, a true negative is occurred when both the prediction outcome and the actual value are n, and false negative is when the prediction outcome is n while the actual value is p. The closer the curve follows the left-hand border and then the top border of the ROC space, the more accurate the test. However, this test is a subjective test, and the results of the test have been determined by an unbiased examiner, well versed with the nuances of neural spike detection, and unaware of the measures used to determine the tests.

	Before Correlation	After Correlation
Day 2	1.85	2.4
Day 25	2.93	3.0
Day 27	3.16	3.28

Table 4.1. SNR improvement when using Correlation

4.3 Evaluative Comparison of Different Techniques

The six different algorithms listed below are all implemented on "Rat 47" and "Rat 46" on different days, with different rates of firing, and with the neural interface degrading over time. A legend has been included on the top of the plot. The objective measures have been printed on the plot 4.3 and 4.3 for easy reference.

The best performing algorithms, in terms of the subjective measures alone, are the algorithms involving PCA, which regularly outperform the algorithms that do not involve PCA. Correlation in general seems to improve thresholding, and a rethresholding further improves the results after a pure correlation. On days with high mastication, PCA obtains a good spike shape, but it gets rid of fewer mastication like signals, and hence it has a better mean spike than all other waveforms. On such days, waveforms with the correlation algorithm show worse mean spikes than waveforms without that algorithm (pure thresholding, and thresholding followed by PCA).

4.4 Results showing Improvement using Intra-Electrode Correlation

The algorithm described above has been implemented on "Rat 47" on different days, with different rates of firing, and with the neural interface degrading over time. The top plot shows the performance when the interface is at its best performance, and the interface gradually degrades with time. The results shown are for Days 1, 11 and 27 of the neural study. The objective measures have been printed on the plot 4.4 for easy reference.

An improvement in SNRs is tabulated below:

Increase in SNR in cases where the number of spikes eliminated by correlation is minimal. Reduction in SNR in cases wherein the number of spikes eliminated

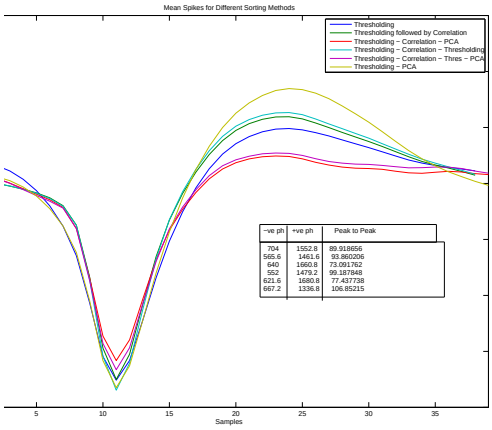

(a) Early Day: Low mastication, high SNR, good neural interface

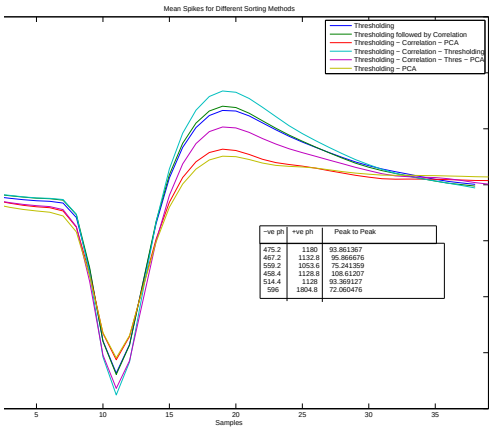

(b) Middle Day 25: Intermediate mastication, SNR, neural interface

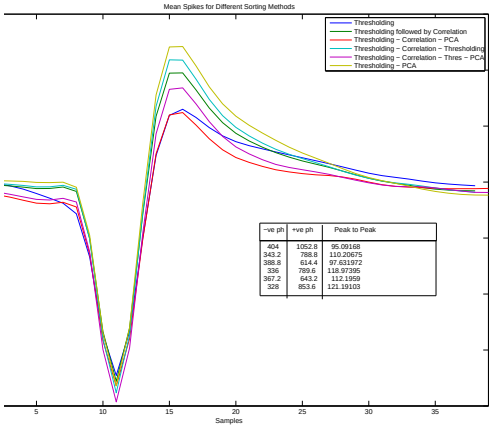

(c) Late Day 27, low SNR, large amount of Mastication present

Figure 4.1. Performance of all algorithms on Rat 47

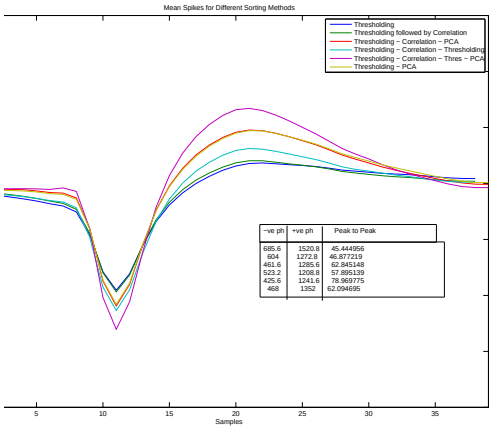

(a) Early Day: Low mastication, high SNR, good neural interface

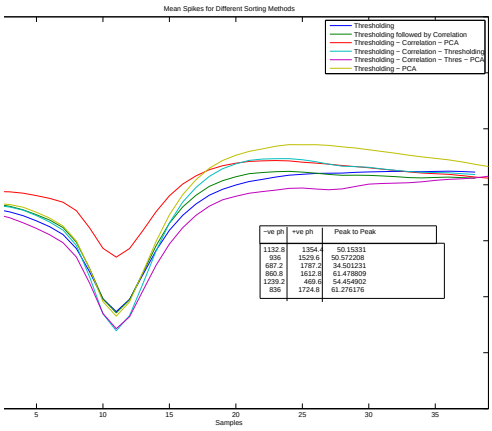

(b) Middle Day 25: Intermediate mastication, SNR, neural interface

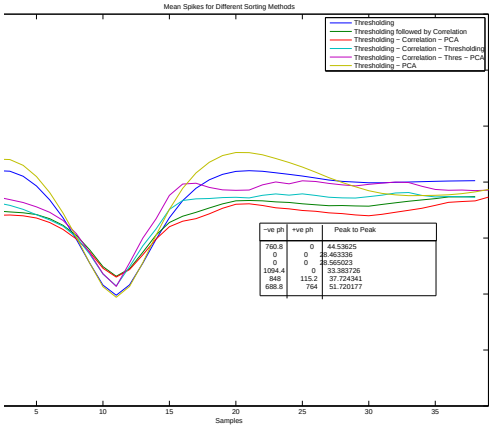

(c) Late Day 27, low SNR, large amount of Mastication present

Figure 4.2. Performance of all algorithms on Rat 46

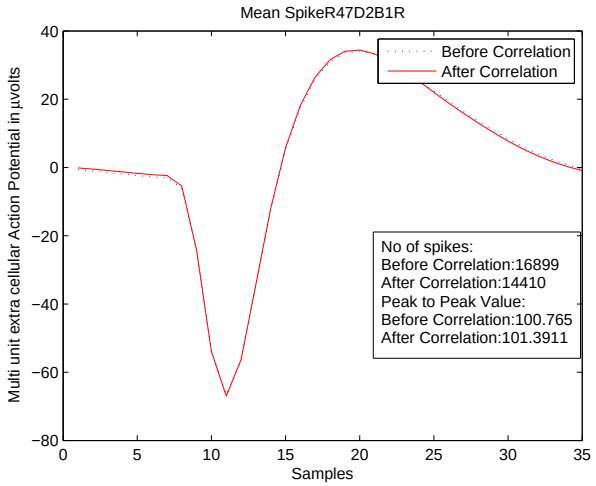

(a) Early Day: Low mastication, high SNR, good neural interface

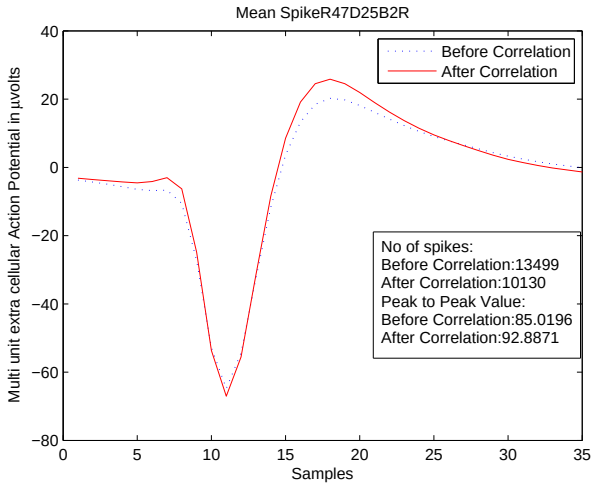

(b) Middle Day 25: Intermediate mastication, SNR, neural interface

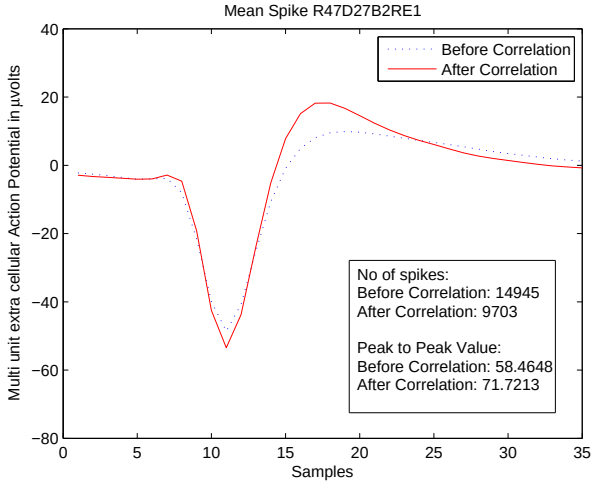

(c) Late Day 27, low SNR, large amount of Mastication present

Figure 4.3. Improvement shown in Mean Spikes using Intra-electrode Correlation

by correlation is significant, thereby pointing to the contribution of non-neuronal signals in the calculation of signal amplitudes.

The number of spikes recorded on the micro-wire arrays was analyzed as well. The number of uncorrelated spikes was comparable in one animal (R46) and the difference was statistically insignificant (p ¡ 0.81). R47 (p ¡ 0.05) and R50 (p ¡ 0.0001) had statistically significant differences in the number of recorded uncorrelated spikes. In both cases the collagenase treated sites recorded higher number of spikes as compared to the control sites. Figure tracks the average number of spikes recorded on the micro-wire array for a period of 1 month.

The objective measures clearly show an improvement in the mean spike shape, and a general increase in sharpness of the waveforms. Also, the number of spikes removed by Correlation increases with an increase in the extent of the mastication levels observed while recording. This is more evident when the receiver operating characteristics are plotted in the graphs in 4.4. Maximum improvement in shape of spike in cases wherein maximum reduction/change in number of detected spike post-correlation.

4.5 Improvements in the PCA algorithm when using Correlation

PCA is a single electrode approach that looks only at the data from one electrode to make a decision. This approach has been clubbed with correlation to improve the ability of PCA to detect mastication like signals. While this is not immediately apparent when generating objective measures only, it becomes completely evident while performing a look test and generating ROCs.

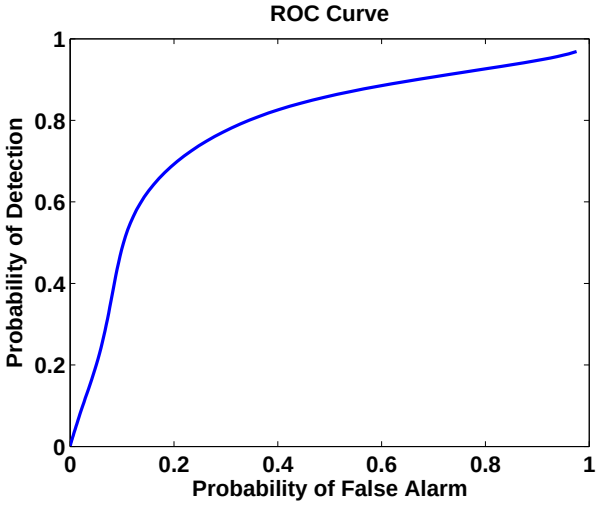

(a) Early Day: Low mastication, high SNR, good neural interface

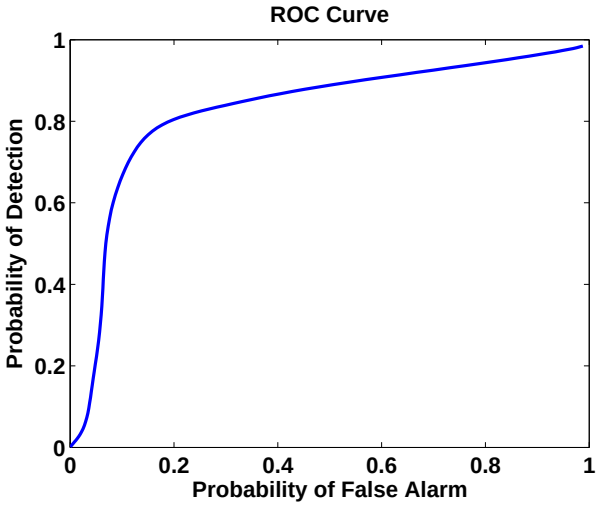

(b) Middle Day 25: Intermediate mastication, SNR, neural interface

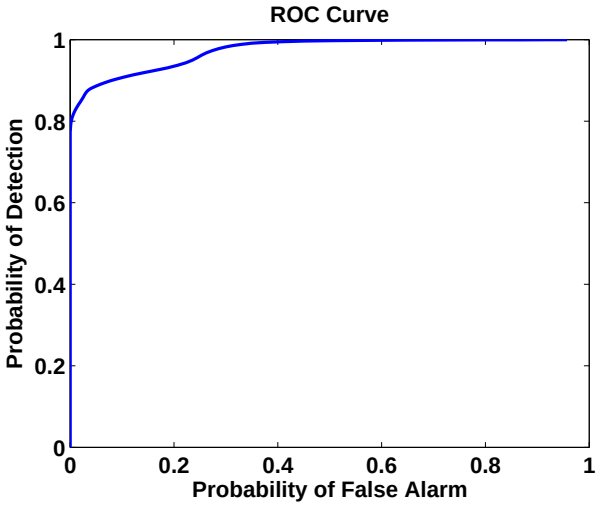

(c) Late Day 27, low SNR, large amount of Mastication present

Figure 4.4. Receiver Operating Characteristics evaluating the Intra-electrode Correlation method

CHAPTER

5

CONCLUSIONS

This thesis shows the implementations of several algorithms for spike detection, and it describes the development of a new scheme for spike detection based on multiple electrode observations. Methods involve those that can be implemented in an online, real time scenario, with minimum amount of training or calibration.

The significant findings of this study can be summarized as follows:

1. A simple thresholding scheme is designed that overcomes the problem of multiple and overlapping spikes. Results are presented showing the efficacy of this scheme, with respect to objective and subjective measures.

2. A multi-electrode method is designed to detect spikes, while looking at data from multiple electrodes that are located close to each other. Mastication waveforms that are very closely resembling spikes are eliminated, thus reducing the probability of false alarm.

3. The different methods are implemented on a large number of data sets, with varying degrees of mastication, Signal to Noise Ratios, and with different

performances of the neural spike interface.

4. PCA performs very well on a dataset with low mastication, and low intra-electrode correlation. However, in an adverse environment, its performance is only somewhat better than a coin-toss.

5. Intra-electrode correlation is a very good metric to detect spikes in a data set where mastication is a problem. It outperforms all other methods, giving a probability of detection of $\geq 90\%$ and an excellent ROC curve.

6. Using this correlation algorithm in conjunction with PCA gives the best performance in terms of objective measures defined, and subjective measures like the ROC curve. The ROC curve of a correlation algorithm with PCA is much better than the ROC curve of a standalone PCA algorithm.

Recommended directions for future studies are:

1. Development of techniques using wavelets, non-linear energy operator etc. for spike detection.

2. Decision fusion of different techniques when working together for the task of spike sorting.

3. Implementation of correlation along with different algorithms that operate only on a single electrode.

4. Development of other methods using multiple electrodes. This includes development of a reference electrode that is an "average electrode" and subtracting this signal from all other electrodes, to remove effects of mastication.

BIBLIOGRAPHY

[1] Spelman, F. A. (1999) "The past, present, and future of cochlear prostheses," *IEEE Eng. Med. Biol. Mag*, **18**, pp. 27–33.

[2] Lui, W. (2002) "Retinal implant: Bridging engineering and medicine," *Proc. IEEE Int. Electron Devices Meeting*, **1**, pp. 492–495.

[3] P. Limousin, P. P. A. B. C. A. D. H., P. Krack and A. Benabid (1998) "Electrical stimulation of the subthalamic nucleus in advanced Parkinsons disease," *New Eng. J. Med*, **339**, pp. 1105–1111.

[4] J. K. Chapin, R. S. M., K. A. Moxon and M. A. Nicolelis (1999) "Real-time control of a robot arm using simultaneously recorded neurons in the motor cortex," *Nature Neurosci*, **2**, pp. 664–670.

[5] Strumwasser, F. (1958) "Long-term recording from single neurons in brain of unrestrained mammals," *Science*, **127**, pp. 469–470.

[6] Nordhausen, M. E. M. . N. R. A., C. T. (1996) "Single unit recording capabilities of a 100 microelectrode array," *Brain Res.*, **726**, pp. 129–140.

[7] Hoogerwerf, K. D., A. C. & Wise (1994) "A three-dimensional microelectrode array for chronic neural recording," *IEEE Trans Biomed Eng*, **41**, pp. 1136–41.

[8] T. Stieglitz, M. S. and K. P. Koch (2005) "Implantable biomedical microsystems for neural prostheses," *Engineering in Medicine and Biology Magazine, IEEE*, **24**, pp. 58–65.

[9] et al., P. R. (2001) "Flexible Polyimide-Based Intracortical Electrode Arrays with Bioactive Capability," *IEEE Transactions of Biomed*, **48**.

[10] Lewicki, M. (1998) "A review of methods for spike sorting: the detection and classification of neural action potentials." *Network: Computation in Neural Systems*, **9**, pp. 53–78.

[11] Musial PG, G. G. K. E. K. J., Baker SN (2002) "Signal to noise ratio improvement in multiple electrode recordings," *Journal of Neuroscience Methods*, **1**, pp. 29–43.

[12] H.D.Brunk (1975) "An Introduction to Mathemeatical Statistics, 3/e," *Xerox College Publishing, MA*.

[13] R. L. Rennaker, A. M. R., S. Street and A. M. Sloan (2005) "A comparison of chronic multi-channel cortical implantation techniques: manual versus mechanical insertion," *Journal of Neuroscience Methods*, **142**, pp. 169–176.

[14] A.K.Jain, R.P.W.Duin, and J.Mao (2000) "Statistical Pattern Recognition: A Review," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, **22**(1), pp. 4–37.

[15] Mallat, S. (1998) *A Wavelet Tour of Signal Processing*, 2 ed., Academic Press, Boston, MA.

[16] Lee, T.-W. and M. S. Lewicki (2002) "Unsupervised classification, segmentation and enhancement of images using ICA mixture models," *IEEE Trans. Image Proc.*, **11**(3), pp. 270–279.

[17] D.L.Donoho (1995) "De-noising by Soft-Thresholding," *IEEE Transactions on Information Theory*, **41**(3), pp. 613–627.

[18] J. Chen, J. F. H., K. D. Wise and J. S. C. Bledsoe (1997) "A multichannel neural probe for selective chemical delivery at the cellular level," *IEEE Transactions on Biomedical Engineering*, **44**, pp. 760–769.

[19] Duda, R. O., P. E. Hart, and D. G. Stork (2001) *Pattern Classification*, John Wiley.

[20] C.J.Veenman, M.J.T.Reinders, E.M.Bolt, and E.Baker (2002) "A Maximum Variance Cluster Algorithm," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, **24**(9), pp. 1273–1280.

[21] R.Duda, P.Hart, and D.Stork (2001) "Pattern Classification, 2/e," *John Wiley, New York*, pp. 915–930.

[22] Kerschen, G. and J.-C. Golinval (2002) "Non-linear generalization of principal component analysis: From a global to a local approach," *Journal of Sound and Vibration*, **254**(5), pp. 867–876.