A Robotic Neural Interface for Autonomous Positioning of Extracellular Recording Electrodes

Michael T. Wolf, Jorge G. Cham, Edward A. Branchaud, Grant H. Mulliken, Joel W. Burdick and Richard A. Andersen

The International Journal of Robotics Research 2009; 28; 1240 originally published online Jun 8, 2009;
DOI: 10.1177/0278364908103788

A Robotic Neural Interface for Autonomous Positioning of Extracellular Recording Electrodes

Michael T. Wolf
Jorge G. Cham
Edward A. Branchaud
Grant H. Mulliken
Joel W. Burdick

Department of Mechanical Engineering,
California Institute of Technology,
1200 East California Boulevard,
Pasadena,
CA 91106,
USA
wolf@robotics.caltech.edu

Richard A. Andersen

Division of Biology
California Institute of Technology,
1200 East California Boulevard,
Pasadena,
CA 91106,
USA

Abstract

In this paper we describe a set of algorithms and a novel miniature device that together can autonomously position electrodes in neural tissue to obtain high-quality extracellular recordings. This robotic system moves each electrode to detect the signals of individual neurons, optimize the signal quality of a target neuron, and then maintain this signal over time. Such neuronal signals provide the key inputs for emerging neuroprosthetic medical devices and serve as the foundation of basic neuroscientific and medical research. Experimental results from extensive use of the robotic electrodes in macaque parietal cortex are presented to validate the method and to quantify its effectiveness.

KEY WORDS—neural interface, neuroprosthetics, brain-machine interface, extracellular recording, neurorobotics, electrode microdrive

1. Introduction

In this paper we present a robotic system that can autonomously position electrodes in cortical tissue to find, optimize, and track the extracellular signals generated by individual neurons. Such signals provide control inputs for emerging neuroprosthetic[1] medical devices that promise to aid the severely handicapped (Wessberg et al. 2000; Serruya et al. 2002; Taylor et al. 2002; Andersen et al. 2004) and are also the targets of a large body of electrophysiological experiments aimed at basic understanding of the nervous system.

The goal of the extracellular recording process (more fully described in Section 2) is to detect and localize in time the occurrence of a neuron's electrical impulses, termed *action potentials* or *spikes*, which are the basis for neural communication and information processing. It is widely accepted that the information output of a neuron is encoded not in the shape of

1. A *neuroprosthesis* is a brain–machine interface that enables a human, via the use of surgically implanted electrode arrays and associated computer decoding algorithms, to control external electromechanical devices by pure thought alone. In this manner, some useful functions that have been lost through disease or accident can be partially restored. Patients who might benefit from such a prosthesis would include those with late stage amyotrophic lateral sclerosis (ALS), severe spinal cord lesions or trauma, and stroke to the motor pathways.

The International Journal of Robotics Research
Vol. 28, No. 9, September 2009, pp. 1240–1256
DOI: 10.1177/0278364908103788

Figures 1–3, 5–11 appear in color online: http://ijr.sagepub.com

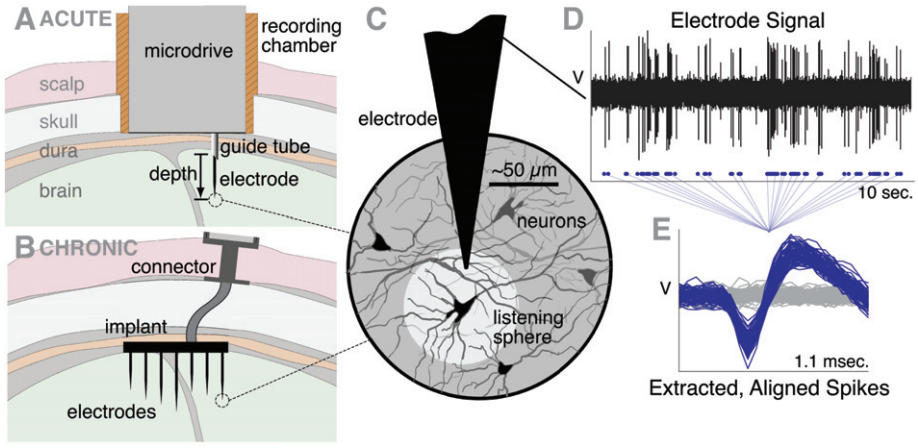

Fig. 1. Extracellular recording environment and example signals: cross-sectional diagrams of (A) acute and (B) chronic recording setups; (C) detail drawing of recording site at electrode tip; (D) 10-second filtered signal sample from an electrode, with (E) the action potential (spike) waveforms extracted from the recording and aligned by their minimum.

its spikes, which are highly stereotyped (see, e.g., Figure 1E in Section 2), but in their timing. A successful extracellular recording, then, is one in which the firing of spikes of an individual neuron can be reliably detected above background noise and distinguished from signals arising from nearby neurons; the neuron is then considered as "isolated". The timing of these spikes may then be analyzed for scientific studies or for control of a neuroprosthesis, decoding the intentions of a paralyzed user.

Whether a cell is successfully isolated, however, relies almost entirely on the effective placement of the uninsulated electrode tip with respect to that cell body. The robotic system described below can autonomously position electrodes so as to initially optimize and then maintain the quality of the recorded signal over long periods of time. Our approach and contributions consists of two parts. First, we present a hierarchical control algorithm that can determine appropriate electrode movement commands to optimize the quality of the recorded signal. This control system can be used with a wide variety of different hardware configurations. Second, we present a novel miniature robotic electrode microdrive[2], which takes advantage of the proposed control algorithms. The robotic electrode paradigm described in this paper can increase the quality and efficiency of neuroscientific research techniques by eliminating the tedious manual process by which electrophysiologists have traditionally optimized electrode placement. In addition, future miniaturized versions of such a system may overcome some of the difficulties inherent in establishing effective, long-lasting neural interfaces that are required for practical neuroprostheses.

Section 2 provides additional further background on the techniques and challenges of extracellular recording in order to motivate the potential contributions of autonomous electrode positioning. Section 3 describes the hierarchical control algorithm that governs the electrode positioning process. Not only can this algorithm be used with the novel robotic microdrive described in Section 4, but it can also be employed to automate or "robotize" existing commercial microdrives. Section 5 provides experimental results obtained by applying the control system, in tandem with the robotic microdrive of Section 4 as well as commercial microdrives, to autonomously record extracellular signals in macaque parietal cortex.

While microdrives have long been used for basic research in neurobiology, our robotic microdrive is novel in its combination of size and autonomy. Few attempts have been made to automate the delicate process of electrode positioning. Fee (2000) has previously demonstrated a method to stabilize intracellular recording electrodes for a period of a few minutes. Also, Baker et al. (1999) have demonstrated a control architecture for an acute microdrive that autonomously advances electrodes until target cells are detected, at which point a human operator must optimize the recorded signal.

A first generation of our robotic microdrive prototype was previously reported by Branchaud et al. (2005) and Cham et al. (2005). The advances in the robot described in this paper provide substantial improvements in terms of signal quality, robustness to biological environments, experimental ease of use, and manufacturability. A preliminary variant of the current microdrive was presented by Cham et al. (2006). An early version of our control system can be found in Branchaud et al. (2006); Nenadic and Burdick (2006). We have developed a significantly more sophisticated electrode positioning algorithm, capable of coping with the many challenges of realistic

2. A *microdrive* is an electromechanical device that can position an electrode along a linear track with micrometer-scale precision. The device itself may be quite large.

recording environments and able to track the signals of a given neuron with greater reliability. Moreover, efforts are underway to develop miniaturized implantable versions of this technology based on micro-fabrication techniques (Pang et al. 2006); the progress reported here supports the development of these devices. Finally, we present the first quantitative analysis of the performance of our approach from many days of use, having logged over a thousand hours of autonomous electrode control in multiple animal subjects.

2. Background: The Extracellular Recording Process

Extracellular recordings of action potentials are made by inserting electrodes into neural tissue. The electrodes are typically sharpened metal wires insulated along their length and exposed at the tip[3]. The tip of a recording electrode must generally lie within about 50 μm of the neuron's soma to provide a useful signal that can be discriminated above the background noise (defining a "listening sphere" around the neuron; see Figure 1C), and a closer proximity may be required to sufficiently distinguish the signals of different neurons (Gray et al. 1995). As summarized below, extracellular recordings can be carried out in an *acute* or in a *chronic* manner. The autonomous electrode positioning algorithm described in this paper can benefit both types of extracellular recordings.

2.1. Acute Recordings

In acute recordings, which are primarily used for scientific research, electrodes are inserted and removed from the neural tissue during each recording session (which typically lasts a few hours). To enable these recordings in cortex, a portion of the skull over the brain region of interest is typically removed and replaced with a sealable cranial *recording chamber* (see Figure 1A); for example, a 16-mm inner-diameter cylindrical recording chamber is a standard used in the neuroscience community. During an acute recording session, a microdrive[4], affixed to the opened chamber, is used to lower one or more electrodes into cortical tissue and then subsequently finely position the electrodes. Electrodes are advanced through neural tissue along a straight line, with the position of each electrode described by its *depth* along this linear track.

While the electrode movement is typically motorized, the electrode's motion is at present manually determined by the experimenter. The process of determining the exact position of each electrode is commonly guided by the use of visual (oscilloscope) and auditory (loudspeaker) representations of the voltage signal, and the experimenter relies on experience and intuition to determine proper electrode placement. The electrode must be placed where signals of disparate neurons are distinguishable (Lewicki 1998) and close enough to a neuron for a high-quality recording, yet far enough away to avoid damaging it. During the course of a typical experiment, the experimenter must monitor the electrode and often reposition it, as the signal can change due to tissue decompression effects (Emondi et al. 2004). The process of isolating and maintaining high-quality neuronal signals thus consumes a significant amount of the experimenter's time and focus.

Simultaneous recordings with many electrodes are becoming an increasingly important technique for understanding how local networks of neurons process information, as well as how computations are coordinated across multiple brain areas (Buzsaki 2004). Commercial microdrives with 16 or more electrodes are now available (Baker et al. 1999). As the number of electrodes increases, the manual task of positioning each electrode to maintain a high-quality neuronal signal becomes intractable for a single experimenter. Data collection in these experiments is essentially limited by how many channels an experimenter can effectively monitor; in our experience, about three or four electrodes is the maximum that can be juggled effectively by an experienced electrophysiologist. Thus, by continually monitoring the signal and automating the process of placing and repositioning electrodes, our system can significantly improve the efficiency and quality of acute multi-electrode studies.

2.2. Chronic Recordings

In chronic recordings, multi-electrode assemblies, which typically consist of fixed geometry bundles of thin wires or arrays of silicon probes, are surgically implanted in the region of interest (Rousche and Normann 1998; Williams et al. 1999; Porada et al. 2000) (see Figure 1B) and remain in place for weeks, months, or possibly years at a time. Such chronic implants enable longer-term scientific studies of larger populations of neurons or can be used as the front end of a neuroprosthesis.

Current chronic recording technology suffers from a number of limitations. The implant's yield (the percentage of the array's electrodes that can record a useful signal) depends largely upon the luck of the initial surgical placement, as many electrodes will likely fall outside the "listening sphere" of any active neuron. Moreover, blood pressure variations, breathing, and mechanical shocks can cause migration of the electrodes in the tissue, leading to further signal degradation (Avezaat and van Eijndhoven 1986; Fee 2000). Finally, reactive gliosis can encapsulate the electrode, diminishing signal quality over time (Turner et al. 1999). All of these effects conspire to limit the usefulness and practical longevity of chronically implanted electrode arrays.

3. Silicon shafts with electrically active recording sites along their shanks may also be used.

4. For example, commercial microdrives from Thomas Recording GmbH, Germany, FHC Inc., USA, Narishige Inc., Japan, NAN Instruments Ltd, Israel, etc. A photo of several microdrives is shown in Figure 5.

A chronic array whose electrodes can be continually repositioned after implantation can overcome some of these limitations. With such an implant, the overall signal yield can be improved by moving the electrodes to optimal neuronal recording sites, even seeking out neurons whose activity is well correlated with the objectives of the neuroprosthesis. While some implantable microdrives have been developed (deCharms et al. 1999; Venkatachalam et al. 1999; Vos et al. 1999; Fee and Leonardo 2001; Kralik et al. 2001), these devices require manual repositioning of the electrodes, such as by turning lead screws. Muthuswamy et al. (2005) have demonstrated a prototype actuated electrode in an acute rat experiment; however, it is unclear whether the high power consumption and limited actuator range of their device will be appropriate for chronic use in primates. Also, an accompanying control algorithm such as ours would still be necessary, as it is not practical to require constant human supervision to adjust the electrodes to achieve optimal signals.

Although the experimental results in this paper focus on acute recordings, our algorithms and demonstrations provide a foundation for future chronic "smart" implantable multi-electrode systems. Chronic tests of autonomous electrode positioning will require new miniaturized actuator technologies: Pang et al. (2006) documents initial attempts at developing miniaturized, biocompatible, actuated electrodes that would enable an implantable device. The microdrive described in Section 4 serves as a testbed for the autonomous electrode approach that will be extended to longer time frames in future trials.

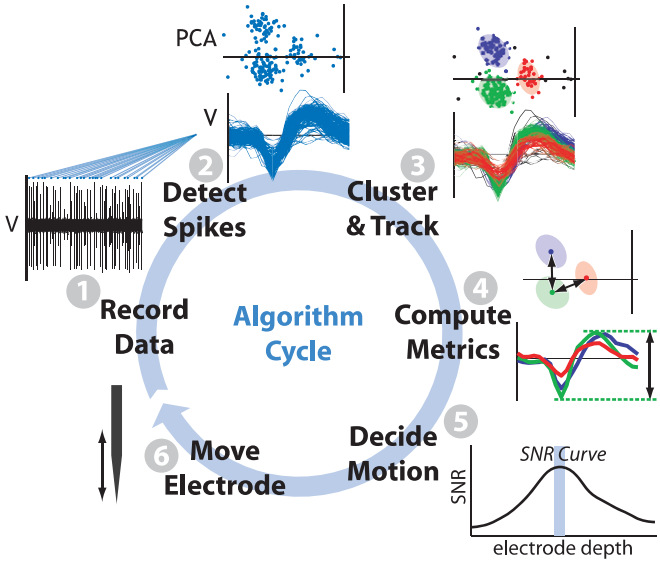

Fig. 2. Diagram of the control algorithm cycle, depicting the key data involved in each step of the cycle. During the kth cycle: (1) a short data sample is recorded during interval T_k, from which (2) spike waveforms are detected, extracted, and aligned. (3) Using their PCA representations, these waveforms are clustered by their generating neurons and associated with the neurons recorded on the previous cycle. (4) SNR and IQM metrics are computed and then (5) used to determine the electrode motion commands to optimize the SNR curve. (6) The electrode is moved to its commanded position.

3. Autonomous Electrode Positioning Control Algorithm

Our algorithm utilizes a hierarchical closed-loop approach to determine, based on the recorded signal and the electrode's position history, the best depth for each electrode. The goal is to place each electrode so that the spikes from an *isolated* neuron can be unambiguously detected in the noisy voltage recording and discriminated from the signals of other nearby neurons. We first review the overall control system structure and then describe its individual components. To make this paper self-contained, we summarize components that have been developed in previous work, while more thoroughly highlighting new contributions. As each electrode is moved independently, we focus on the processing steps for a single electrode; these steps are run in parallel for each electrode in a multi-electrode microdrive.

The control algorithm operates on a cycle (see Figure 2). Let these cycles be indexed by the integer k, $k = 1, 2, \ldots$. The cycle begins with recording the electrode signal over a short sampling interval (denoted by T_k, which is typically of duration 10–20 seconds) while the electrode is stationary, followed by analysis of this signal to determine whether and how the electrode should be repositioned, and ending with the movement of the electrode to a new position (if necessary).

A hierarchical control algorithm determines the electrode movement commands. The inner-most loop of this algorithm (Section 3.3) attempts to isolate an individual neuron by optimizing the quality of the recorded signal via small local movements of the electrode tip, assuming that the tip is close enough to a neuron for the isolation process to be possible. The outer control structure consists of a finite state machine supervisory controller (Section 3.4) which has several purposes. First, it manages the neuron isolation process: it moves the electrode until a region of sufficiently strong neuronal signal sources is found and then further searches this region to acquire the information needed to initiate the isolation procedure. In addition, the supervisory system (which is a novel contribution of this paper) handles several of the complicating realities of the extracellular recording process. Of course, to provide the neuronal signal metrics that are needed by these control algorithms, several signal processing steps are required to identify and sort the spike waveforms from the raw electrode signal (Section 3.1). Most of these preprocessing tasks have traditionally been performed manually in electrophysiology experiments; producing automated and unsupervised methods

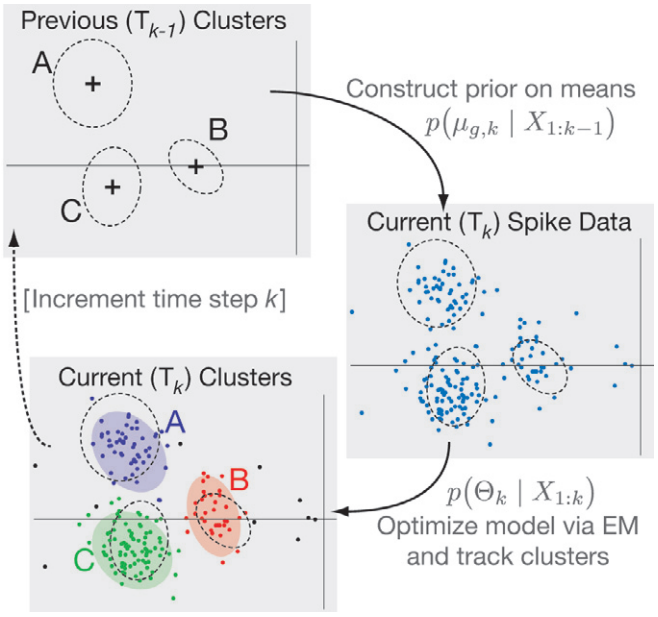

Fig. 3. The clustering process. Clusters identified on the previous recording interval T_{k-1} are used to construct priors on the locations (i.e. means) of clusters on the current recording interval T_k. The current data are clustered via maximum *a posteriori* (MAP) optimization of a Gaussian mixture model via expectation–maximization (EM), and the new clusters are matched to previous clusters to "track" neurons.

presents significant challenges in their own right and has been a contribution of this research program.

3.1. Signal Preprocessing Steps

The first step of unsupervised signal processing on the electrode's recorded signal is *spike detection*, which identifies the action potential events in the raw electrode signal of interval T_k. We employ a wavelet-based method developed by Nenadic and Burdick (2005) specifically designed for this application. By projecting the electrode signal onto a specially designed wavelet basis, spike-like waveforms can be detected in the raw signal, and short intervals ($\sim$1.1 ms in length) of the signal centered on the putative spike occurrence are extracted (see Figure 1D and E). Let $s_{i,k}$ denote the ith spike waveform detected and extracted during the interval T_k. All of the spike-like waveforms found during T_k are temporally aligned by their waveform minima in preparation for the next steps.

A single electrode's signal may contain action potentials generated by multiple neurons lying near the electrode tip, and the detected spikes must be grouped by their generating neurons: a process known as *spike sorting*. Spike sorting is a critical task, as the metrics calculated from the signals of each distinct neuron are vital to both the electrode positioning algorithm, which is trying to maximize signal quality (defined in Section 3.2), and to the scientific or prosthetic uses of the recorded data, which generally rely on neuronal firing rates. If spikes are incorrectly classified, these metrics may be severely corrupted. In addition, we must *track* the signals of particular neurons over sequential recording intervals by matching current signal clusters (identified in T_k) with previous clusters (identified in T_{k-1}).

To more efficiently discriminate between the signals of different neurons, the extracted, aligned spike waveforms are first projected onto a two-dimensional principal component analysis (PCA) basis. We assume that the distribution of each neuron's spikes may be modeled as Gaussian in this feature space (although the center, orientation, and size of each distribution is *a priori* unknown). The use of both a two-dimensional PCA basis and a Gaussian distribution are common practice in spike sorting (Lewicki 1998), although our method accepts any feature space in which the Gaussian assumption may reasonably hold (and it may be adapted for non-Gaussian spike feature distributions).

Let $x_{i,k}$ denote the projection of spike waveform $s_{i,k}$ onto the PCA basis. Our unsupervised spike sorting procedure (Wolf and Burdick 2007) is based on the optimization of a Gaussian mixture model, in which each Gaussian component models the probability that the spike $x_{i,k}$ is generated from a particular neuron, represented by component (or *cluster*)[5] C_g, $g = 1, \ldots, G_k$. Assuming that the measurements of the spike waveforms detected during T_k are independent, the likelihood function for the ensemble of waveforms detected in T_k, $X_k = \{x_{i,k}\}_{i=1}^{N}$, takes the form

$$p\left(X_k \mid \Theta_k\right) = \prod_{i=1}^{N} \sum_{g=1}^{G_k} \pi_{g,k} f_{\mathcal{N}}\left(x_{i,k} \mid \mu_{g,k}, \Sigma_{g,k}\right), \quad (1)$$

where $\Theta_k = \{\pi_{g,k}, \mu_{g,k}, \Sigma_{g,k}\}_{g=1}^{G_k}$ represents the mixture model's parameters[6]. The multivariate Gaussian distribution $f_{\mathcal{N}}(x_{i,k} \mid \mu_{g,k}, \Sigma_{g,k})$ of the spikes associated to the gth generating neuron during T_k is defined by its mean $\mu_{g,k}$ and covariance matrix $\Sigma_{g,k}$. The mixture weight $\pi_{g,k}$ of component C_g describes the probability that the observed spike is generated by the gth neuron.

Traditionally, the model parameters of the Gaussian mixture are found by optimizing the likelihood (1) via expectation–maximization (EM) (Fraley and Raftery 1998). As a byproduct, the EM solution associates spikes to a specific

5. The total number of generating neurons active in interval T_k, G_k, is assumed to be unknown. We implement the common practice of fitting models for a range of G_k (typically up to $G_k = 5$) and selecting the best resulting model (in a way that penalizes over-fitting).

6. In practice, we also include an additional "background" mixture component of uniform density to capture "outliers": detected events that do not fit well into the Gaussian components.

neuron, i.e., the spikes are "clustered". This method has proven effective for many spike sorting applications (Lewicki 1998). However, when the recording application involves repeated sampling and clustering over time, our experience has shown that the inconsistency of the output of conventional clustering methods prevents accurate tracking of the neurons' identities across sampling intervals. If a neuron cannot be tracked, it is very difficult to predict how electrode movement will affect its recorded signal quality. Thus, producing more consistent, accurate clusters is essential for downstream processing in our control system. Consequently, we employ a Bayesian clustering approach, which optimizes the posterior density

$$p\left(\Theta_k \mid X_{1:k}\right) \propto p\left(X_k \mid \Theta_k\right) p\left(\Theta_k \mid X_{1:k-1}\right)$$

instead of the likelihood (1), where $X_{1:k}$ denotes all data recorded during cycles 1 through k. As the cluster of signals associated to the gth neuron in interval T_k is apt to be similar in location to the signal cluster generated by the same neuron in interval T_{k-1}, an informative prior is placed on the means $\mu_{g,k}$ of the Gaussian mixture components found in T_k based on the previous time step's $\hat{G}^{k-1}$ cluster locations:

$$p\left(\Theta_k \mid X_{1:k-1}\right) \propto \prod_{g=1}^{G} p\left(\mu_{g,k} \mid X_{1:k-1}\right)$$

$$= \prod_{g=1}^{G}\left(\sum_{j=1}^{\hat{G}^{k-1}} \frac{1}{\hat{G}^{k-1}} f_{\mathcal{N}}\left(\mu_{g,k} \mid \hat{\mu}_{j,k-1}, S_{j,k|k-1}\right)\right),$$

where $S_{j,k|k-1}$ is the covariance matrix associated with the prediction that the mean $\mu_{g,k}$ will lie at location $\hat{\mu}_{j,k-1}$. This maximum *a posteriori* (MAP) approach also associates clusters found in T_k to the clusters found in T_{k-1}. In addition, we use the clusters found in T_{k-1} as the "seed clusters" needed to initialize the EM algorithm and the model selection process in interval T_k. For more details on these techniques, including methods to properly account for newly appearing and disappearing neurons in T_k, see Wolf and Burdick (2007). This method provides significantly more consistent clusters than traditional methods, while also solving the data association problem of spike clusters across recording intervals, thus effectively tracking the signals of persisting neurons throughout the recording session.

3.2. Signal Quality and Isolation Metrics

From the preprocessed data, two signal metrics are calculated in interval T_k.

- A *signal quality metric* (SQM) determines the general quality of the extracellular signals associated with a particular neuron.
- An *isolation quality metric* (IQM) measures the separation of one neuron's waveforms from those of other neuronal signals that appear in the same recording interval.

The SQM is the algorithm's main target, and the *dominant* neuron is chosen as that whose signals have the highest average SQM over a recent time period. This is the neuron whose signal is to be optimized by the electrode's movements. Although other choices of SQM are possible (see Branchaud (2006) for examples), we hereafter use the signal-to-noise ratio (SNR) as the signal quality metric. In this application, it is defined as the mean peak-to-peak amplitude of the neuron's waveforms detected in T_k divided by the RMS amplitude of a spike-free noise sample taken during interval T_k.

As a neuronal signal is only valuable if it can be distinguished from those of nearby neurons, the IQM measures the "isolation" of the dominant neuron's waveforms from other detected spikes. The IQM is based on the *isolation distance* (ID) (Harris et al. 2001), which, for cluster C_g containing N_g spike samples, is defined as the Mahalanobis distance between its center μ_g and the N_gth closest spike not in cluster C_g (denoted by x_j):

$$\mathrm{ID}_g = \sqrt{(x_j - \mu_g)^{\mathrm{T}} \Sigma_g^{-1} (x_j - \mu_g)}.$$

The ID is the radius of the smallest ellipse (with shape defined by Σ_g) containing all of the spikes in cluster C_g and an equal number of spikes not in cluster C_g (in effect, a measure of the "moat" around cluster C_g). Note that the SQM is calculated from the spike waveforms, while the IQM is computed in feature space (PCA basis).

3.3. The Isolation Control Loop

Based on the processed neural data and the quality metrics just defined, the isolation control loop determines whether repositioning the electrode can improve the dominant neuron's signal quality. In the idealized, but unrealistic, scenario where only one neuronal signal source is ever present, the algorithm commands the electrode motion solely to increase the dominant neuron's SNR, as outlined below.

Detailed computational models (Nenadic and Burdick 2006) of the extracellular field generated around a typical cortical pyramidal neuron show that when the electrode tip is within the neuron's "listening sphere", the variation of the neuronal signal's SNR with respect to electrode position traces out a unimodal curve which we dub the *SNR curve* (see Figure 2 (step 5) or the data in Figure 8(b)). Let u denote the position of the electrode tip along its linear track. Let $Y(u)$ denote the SNR curve. The goal is to find the peak of this curve and then maintain the electrode position sufficiently close to this peak. As neural signals are highly noisy, the metric Y should be considered a random variable with an associated regression function $M(u) = E(Y \mid u)$, where $E(\cdot \mid \cdot)$ denotes conditional expectation. This regression function is *a priori* unknown, except

that it has a unimodal shape. Only noisy observations of the SNR, obtained via the preprocessing steps summarized above, are available. In order to optimize the SNR using only noisy samples, the isolation process adaptively estimates the regression function (the smoothed SNR curve), and the electrode's movements are chosen to seek the extremal point of the adaptively evolving SNR curve.

The regression function model $M(u)$ is assumed to be a linear combination of basis functions: $M(u, n_k, B_k) = \sum_{i=1}^{n_k} b_{i,k}\psi_i(u)$, where n_k is the number of basis functions employed during cycle k, and $B_k = [b_{1,k}, b_{2,k}, \ldots, b_{n_k,k}]^{\mathrm{T}}$ are the corresponding expansion coefficients. The model parameters B_k and model complexity n_k must be estimated from SNR observations and adaptively updated as new data become available. For a given model estimate, the electrode's next position, u_{k+1}, is determined as

$$u_{k+1} = u_k + C\,|H_k|^{-1}\,\xi_k, \tag{2}$$

where $C > 0$ is an appropriately chosen scale factor, and ξ_k and H_k are the estimates of the first and second derivatives, respectively, of $M(u)$ at the electrode's current position, u_k. Note that (2) represents a stochastic version of Newton's method. Convergence of the electrode position to the maxima of the SNR curve is considered attained at iteration k^* if $C|H_{k^*}|^{-1}\xi_{k^*} < \epsilon$, where ϵ is a tolerance chosen by the user. The position u_{k^*} is then declared the optimal electrode placement, whereupon the finite state machine supervisory controller transitions to a "maintain" mode (see Section 3.4). The regression function $M(u)$ is estimated as follows.

While many basis function choices are possible, we have found that polynomial bases can sufficiently capture the geometry of unimodal SNR curves (see Nenadic and Burdick (2006)) and greatly simplify the estimation process. For polynomial bases, the regression function after k iterations is

$$\hat{M}(u, n_k, B_k) = \sum_{i=1}^{n_k} b_{i,k}u^{(i-1)}.$$

Let $\{u_1, u_2, \ldots, u_k\}$ be a sequence of (electrode) positions with the corresponding SNR samples denoted by $\mathcal{Y}_{1:k} = \{\mathbf{y}(u_1), \mathbf{y}(u_2), \ldots, \mathbf{y}(u_k)\}$. At each electrode location u_j $(j = 1, 2, \ldots, k)$, multiple observations of SNR have been taken (one for each isolated neuronal waveform), i.e. $\mathbf{y}(u_j) = \left[y_1(u_j), y_2(u_j), \ldots, y_{N_j}(u_j)\right]^{\mathrm{T}}$, where N_j is the total number of observations at u_j (this number may vary across sampling intervals).

Determining the "correct" number of modes, n_k, amounts to a *model selection* problem. Given a family of candidate models $\{\hat{M}(u, n_k, B_k) : n_k = 1, 2, \ldots, N_{\max}\}$, the goal is to select the order of the model that is most probable in view of the data $\mathcal{Y}_{1:k}$ and any prior information, I. The probability of the model $\hat{M}_{n_k}$ given $\mathcal{Y}_{1:k}$ and I follows from Bayes' theorem

$$P(\hat{M}_{n_k} \mid \mathcal{Y}_{1:k}, I) = \frac{p(\mathcal{Y}_{1:k} \mid \hat{M}_{n_k}, I)P(\hat{M}_{n_k} \mid I)}{p(\mathcal{Y}_{1:k} \mid I)},$$

$$n_k = 1, 2, \ldots, N_{\max}, \tag{3}$$

where $\hat{M}_{n_k}$ is short for $\hat{M}(u, n_k, B_k)$ with fixed n_k. Here, I represents the model selection result obtained in the previous interval T_{k-1}: the posterior $P(\hat{M}_{n_{k-1}} \mid \mathcal{Y}_{1:k-1}, I)$ calculated at iteration $k-1$ can be used as the prior at iteration k in (3). The recursion is initialized with a uniform prior density $P(\hat{M}_{n_{k_0}} \mid I) = 1/N_{\max}$, where k_0 denotes the smallest admissible number of iterations, below which there is an insufficient amount of data to reliably model the regression function (see Section 3.4.1). The model order is chosen to maximize the posterior probability (3), i.e.

$$n_k^* = \arg\max_{1 \le n_k \le N_{\max}} P(\hat{M}_{n_k} \mid \mathcal{Y}_{1:k}, I), \quad k = k_0, k_0 + 1, \ldots.$$

To perform this maximization, the posterior $P(\hat{M}_{n_k} \mid \mathcal{Y}_{1:k}, I)$ of each candidate model $\hat{M}_{n_k}$ must be evaluated by marginalizing the unknown parameters B_k. With a Gaussian noise assumption and polynomial bases, the marginalization of B_k can be performed analytically (Nenadic and Burdick 2006).

Once the optimal model order n_k^* at iteration k is known, the parameters of the model $\hat{M}(u, n_k^*, B_k)$ are estimated by a linear least-squares method:

$$B_k^* = \arg\min_{B_k} \left\{ \sum_{j=1}^{k} \|\Psi_{j,k}B_k - \mathbf{y}(u_j)\|^2 \right\},$$

$$k = k_0, k_0 + 1, \ldots,$$

where the matrix $\Psi_{j,k} \in \mathcal{R}^{N_j \times n_k^*}$ consists of N_j identical rows given by $[1, u_j, \ldots, u_j^{(n_k^*-1)}]$. Once the optimal parameters B_k^* are estimated, the optimal model $\hat{M}_k^*(u) \triangleq \hat{M}(u, n_k^*, B_k^*)$ at iteration k is fully specified. From this result the gradient and Hessian of the optimal model are then used in (2) to determine the electrode movement.

As sudden large electrode movements are unacceptable, the maximum step size is limited by a constant $\Delta_{\max}$, chosen by the experimenter. This is especially useful for iterations where the optimal model is found to be a straight line ($n_k^* = 2$), which results in $H_k = 0$ and infinitely large step size in (2). Likewise, if for some $k > k_0$ we obtain $\hat{M}^k(u) = b_{1,k}^*$, i.e. $n_k^* = 1$, then $\xi_k = 0$ and the recursion (2) breaks down. In this case we use a simple control strategy:

$$u_{k+1} = u_k + \Delta_{\text{sample}}, \tag{4}$$

where Δ_{sample} is a constant Section 3.4 covers the choice of step size parameters and the search for admissible gradients in more detail.

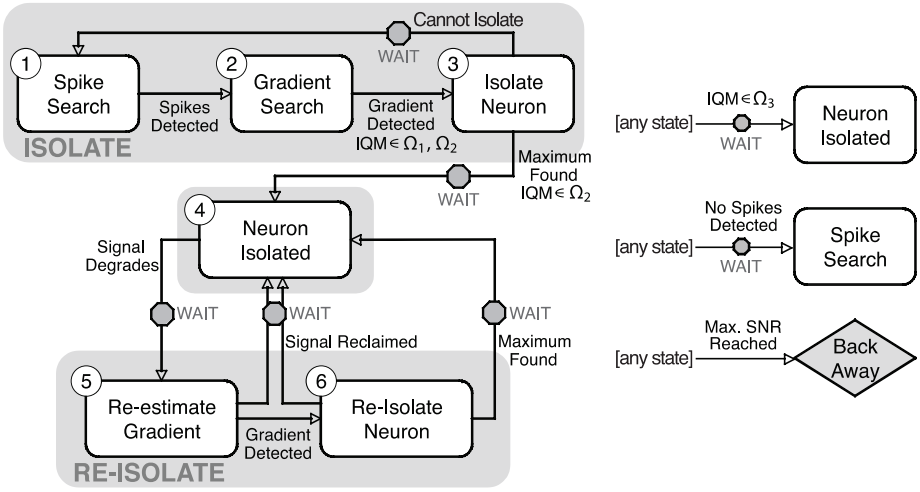

Fig. 4. Diagram of the *supervisory finite state machine* (SFSM) with transition criteria noted. States are grouped into three *modes* (isolate, isolated, and re-isolate) for convenience, as referenced in Section 5.2. Transitions with WAIT must meet transition criteria in R consecutive cycles, reducing sensitivity to transients. Transitions on the right may be made from any state.

To implement this algorithm on an actual hardware device, a low-level electrode positioning feedback control algorithm is needed to execute the commanded change in electrode position from u_k to u_{k+1}. In practice, the low-level controller need not perfectly move the electrode by the distance $\Delta u_k = u_{k+1} - u_k$. The isolation feedback loop described in this section will work as long as the actual electrode displacement proportionally increases the signal quality with each command. While the convergence of the algorithm is reasonably tolerant to errors in electrode positioning, such errors can slow the rate of convergence.

The isolation process described above assumed only one neuron's signal is present during each recording interval T_k ($k = 1, 2, \ldots$). The realistic presence of multiple neuronal signal sources is partially addressed through the spike sorting and neuron tracking procedure of Section 3.1, after which this process uses the signals of just the *dominant* neuron. Additional complications that arise in the presence of multiple signal sources are addressed through the use of a finite state machine for supervisory control, a contribution of this paper.

3.4. Finite State Machine Supervisory Controller

To manage the basic neuron isolation process, while also accounting for many additional challenges of practical extracellular recording, we use a finite state machine architecture to guide the overall electrode movement process. This system is termed the *supervisory finite state machine* (SFSM). During each algorithm cycle, the electrode movement decision depends on the current state of the SFSM, with individual states and state transitions crafted to guide behavior appropriate to seeking and isolating neurons. A prototypical pathway of state transitions is described below to describe the most common issues and SFSM operation, and later sections provide insight into how the SFSM supervisory controller copes with numerous recording challenges.

3.4.1. Prototypical Execution Pathway of the SFSM

When electrodes are first lowered into neural tissue, the electrode tip may not lie in electrically active tissue. The SFSM starts in the *spike search* state (see numbered states in Figure 4), whose goal is to find an electrically active tissue region. In this state, the electrode moves in increments of Δ_{search} ($\sim$20 μm) until a sufficient number of spikes are detected in the interval T_k (according to a minimum firing rate set before the experiment), at which point the SFSM transitions to the *gradient search* state. The *gradient search* state seeks to determine whether a viable SNR curve can be constructed. Observations of the SNR are made at regular intervals of Δ_{sample} ($\sim$10 μm) until k_0 observations are completed (we use $k_0 = 3$ in our experimental studies), at which point the optimization procedure of Section 3.3 determines the most likely order n_k^* that fits the SNR observations. As described above, if $n_k^* = 1$ the electrode continues in steps of Δ_{sample} (the SFSM stays in the *gradient search* state). If $n_k^* > 1$, a potentially viable SNR curve has been found (i.e. a nearby neuron can probably be isolated), the SFSM transitions to the *isolate neuron* state.

As long as the SFSM remains in the *isolate neuron* state, the algorithm described in Section 3.3 operates, updating the SNR curve with new observations and moving the electrode toward the estimated maximum. When the maximum of the SNR curve is reached, the SFSM state transitions to the *neuron*

Table 1. Key Intervals of the IQM.

Interval	Definition	Description
Ω_3	$\gamma_3 \leq$ IQM	Neuron is well isolated; immediately stop and declare isolation as further movement may damage neuron.
Ω_2	$\gamma_2 \leq$ IQM $< \gamma_3$	Neuron is acceptably isolated, if maximum of SNR curve is reached.
Ω_1	$\gamma_1 \leq$ IQM $< \gamma_2$	Isolation quality is high enough to follow gradient, but not high enough for acceptable isolation.
Ω_0	IQM $< \gamma_1$	Isolation quality is too low for reliable measurements; do not follow gradient for stochastic optimization.

isolated state, but only if certain IQM conditions are also met (see Section 3.4.2).

In the *neuron isolated* state, the electrode generally remains stationary while the SNR is continually monitored over successive sampling intervals. Often, the dominant neuron will drift away from the electrode (see Section 3.4.3), causing a decrease in SNR. When the SNR drops below a percentage (typically 85%) of its value at the original isolation, the SFSM transitions to the *re-estimate gradient* state in an attempt to reposition the electrode to maintain the high-quality isolation. The *re-estimate gradient* state moves the electrode in increments of Δ_{resample} ($\sim$5 μm) to find a new gradient. In this state, the electrode is first retracted, as neurons most commonly drift up toward the electrode due to tissue decompression. Once a new gradient is found, a transition is made to the *re-isolate neuron* state, where the optimization procedure is again used to isolate the neuron. If, at any time in the *re-estimate gradient* or *re-isolate neuron* states, the SNR reaches or exceeds the SNR value obtained during the original isolation, the *neuron isolated* state is triggered.

3.4.2. Role of the IQM

The IQM strongly affects the SFSM state transitions. First, to transition from *isolate neuron* to *neuron isolated* state (when the SNR curve maximum is reached), the IQM value must be sufficiently high (above a threshold γ_2) to consider the dominant neuron's signals distinguishable from those of other neurons. Otherwise, the dominant neuron's cell body is likely too far from the electrode's path to provide a good isolation, and thus the SFSM transitions to the *gradient search* state in order to find another suitable neuron. Second, a transition to the *neuron isolated* state will occur from *any* SFSM state when the IQM is very high (above γ_3), even if the SNR curve peak has not been reached. In this case, the neuron probably lies close to or on the electrode's path, and continued advancement to possibly improve the SNR is not worth the risk of damaging the neuron.

Finally, below and IQM threshold γ_1, the isolation of the dominant neuron is of such poor quality that estimates of its SNR cannot be trusted. In such cases, the SFSM transitions to the *gradient search* state, starting a search for a new neuron. In summary, the thresholds described above divide the IQM range into four distinct intervals $\{\Omega_l\}_{l=1}^{3}$, listed in Table 1.

3.4.3. Other Challenges and Related SFSM Transitions

In addition to the challenges described above, several other practical difficulties commonly arise in extracellular recording experiments. Most of these issues are well known to practicing electrophysiologists but have not been systematically reported or studied. The appendix provides interested readers with the details of such challenges and the SFSM components that address them.

4. A Robotic Multi-electrode Microdrive Prototype

This section presents a mesoscale robotic microdrive mechanism that can finely position electrodes in neural tissue using the control algorithm described above, a preliminary version of which was described by Cham et al. (2006). While the control algorithm described above has been used to automate other microdrive mechanisms, the mechanism presented in this section has been specifically designed as a testbed for this approach, and as a means to develop the specifications for future miniaturized implantable devices. The current design is also immediately useful for neuroscience experiments where a small microdrive is desired (e.g., allowing fewer restraints on the subject's head).

4.1. Goals and Challenges

Our design is driven primarily by the desire to use a microdrive in both acute and *semi-chronic* (that is, lasting for days or weeks at a time) experiments, although only acute results are presented in this paper. In order to more easily integrate into current electrophysiological experiments, the microdrive

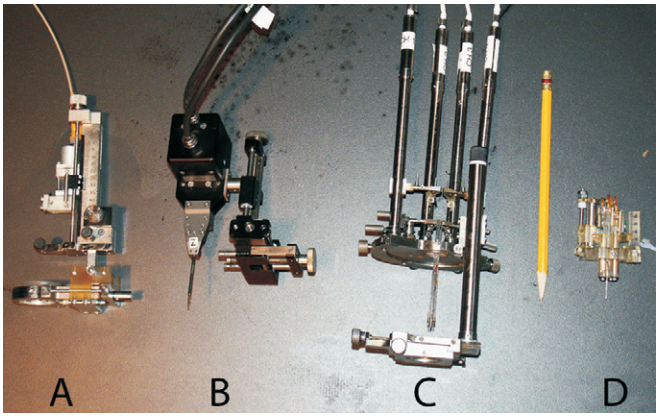

Fig. 5. Photograph of (D) our robotic microdrive prototype alongside three commercial microdrives that have been automated by our control system (see Section 3): (A) FHC; (B) Thomas Recording; (C) NAN Instruments. A pencil (∼18 cm) is included for scale.

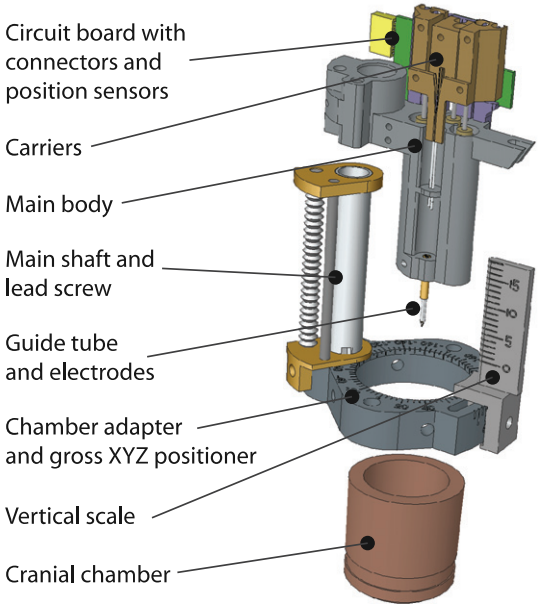

Fig. 6. Exploded view of the electrode microdrive structure.

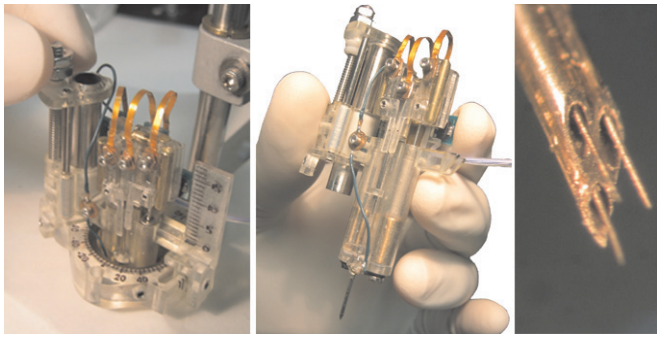

Fig. 7. (a), (b) Photos of the robotic microdrive; (c) detail of the guide tube with electrodes.

should affix to a standard cranial recording chamber used in non-human primate research, but still allow the animal subject free movement and comfort without significant risk of injury. These requirements imply a smaller size and mass than can be obtained from commercial microdrives (see Figure 5). The necessary compactness of the electrical pathways can increase noise and interference in the recorded signal, and the size limitations also restrict the number of actuators, and hence recording electrodes, that can be packaged in the device. The device must be biocompatible and secure against leaks and impacts. The interior of the cranial chamber must be sealed from debris (and the fingers of the animal subject). To enable a variety of research objectives, it must be possible to insert the electrodes over a range of locations within the recording chamber in order to explore multiple brain areas and perhaps to avoid damaging an area by repeatedly piercing the same tissue.

4.2. Design

Figure 6 shows a schematic diagram of the robotic microdrive mechanism. The microdrive's main body encases three piezoelectric linear actuators and provides mountings for an electrode guide tube and a circuit board. The actuators and their associated controllers (Klocke Nanotechnik, Germany) operate on an impulse drive principle that provides both high precision (sub-micrometer position steps) and long range of motion (about 5.6 mm). This range of motion is sufficient to reach the target cortical regions of many acute electrophysiology experiments.

Each linear actuator moves a *carrier*, to which the electrodes are attached both electrically and mechanically. The platinum–iridium electrode wires (Alpha Omega Co., USA) are insulated with glass along their length (except at the recording tip and the back end). Each electrode is loaded tail-first through the guide tube and corresponding carrier, and then screwed to the carrier. The electrode signals are routed to the circuit board (and subsequently by a connector to an off-board amplifier and digitizing computer) via flexible polyimide-coated copper strips. Hall-effect sensors mounted on the circuit board measure electrode depth with 1 μm precision.

The body assembly is held to a recording chamber adapter via a shaft. Rotation of the body assembly around the main shaft axis, combined with rotation of the chamber adapter on the chamber rim, places the guide tubes over any location within a 12-mm diameter circular area inside the chamber. After setting the transverse guide tube position within the chamber, the microdrive is lowered via the vertical screw until the

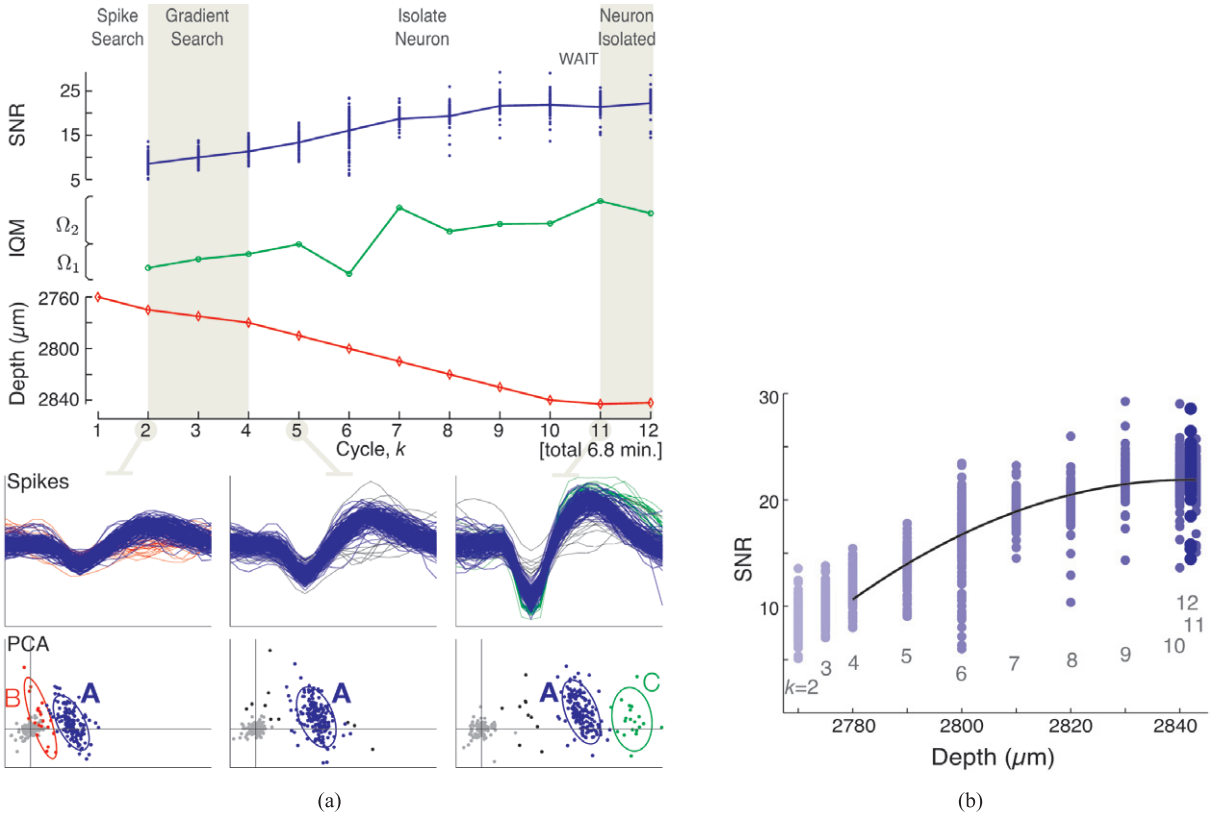

(a) (b)

Fig. 8. Prototypical path to neuron isolation [NAN drive]. (a) The SFSM state, dominant neuron SNR and IQM, and electrode depth over time, with detail on three selected time steps. Each step's detail shows spike waveforms and their PCA representations (using the same axis scales). 2-sigma ellipses in PCA basis designate spike clusters, which are labeled by the generating neuron; the cluster with the bold letter is considered the dominant neuron, for which the SNR and IQM are shown above. (b) SNR curve for the same time period: SNR versus electrode position, with algorithm time step (k) labels.

guide tube pierces the dura. The stainless steel guide tube protects the fragile electrodes (see Figure 7) during this process. Dura penetration is a delicate manual procedure, as advancing the guide tube too far can damage brain tissue. A vertical flange with millimeter rulings aids the experimenter, as well as tactile feedback through the lead screw. Teflon bearings on each shaft reduce friction and provide device stability. Once the guide tube has pierced the dura, the electrodes are deployed with their position determined by the algorithms described above.

4.3. Manufacturing

The microdrive was manufactured using stereolithography (SLA). As processes such as SLA enable complex geometries to be made as one solid piece, this fabrication approach minimized the number of parts needing to be fastened and sealed to assemble the system. Other advantages include the ability to custom-modify each device to fit a particular animal subject. However, the polymers used in standard SLA processes are not biocompatible. To provide biocompatibility, the fabricated parts are coated with a 20 μm layer of Parylene, a US Pharmacy class-VI implantable material that can be deposited using room-temperature low-pressure chemical vapor deposition. The final assembled prototype weighs 26.1 g.

The microdrive presented in this paper offers significant improvements over a first-generation prototype (Cham et al. 2005). The newer prototype's design reduces (by more than half) the time needed for electrode loading and device cleanup. Second, the design increases ease and reliability of dura penetration owing to improved visual and tactile feedback. Third, the use of SLA parts increased robustness to breakage and leakage of biological fluids, and made the microdrive components easier to repair, replace, and also revise. Finally, the new design reduced total device weight by nearly one-half, primarily owing to the elimination of metal parts. Taken together, these solutions to practical challenges have enabled our miniature electrode microdrive to overcome issues of usability and flexibility, which are key barriers to adoption of such technology within the electrophysiology community.

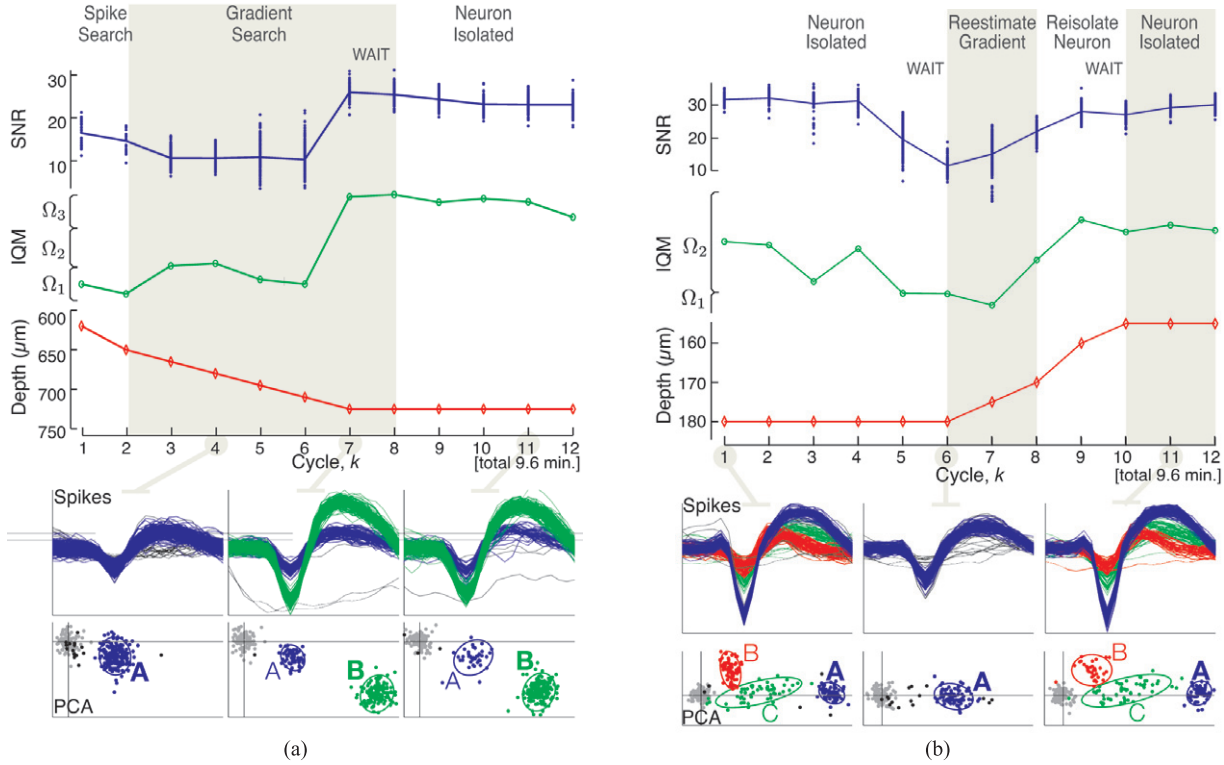

(a) (b)

Fig. 9. Alternative transitions to *neuron isolated* [CIT drive]: (a) isolation by IQM; (b) a re-isolation process. See plot conventions in Figure 8(a).

5. Experimental Results

This section demonstrates our approach in macaque parietal cortex. These data sets were obtained in acute recording sessions, with the microdrive affixed to a standard cranial chamber (see Section 2) and autonomously controlled by our electrode positioning algorithm. Signals, recorded by glass-insulated metal electrodes of approximately 1.5 MΩ impedance at 1 kHz, were amplified and filtered (Plexon, Inc.) and then interfaced to the controlling computer via an analog-to-digital data acquisition card (National Instruments).

As noted earlier, our algorithm has been designed to interface with a variety of electrode microdrives, essentially robotizing these devices. We include data obtained from a microdrive manufactured by NAN Instruments as well as from our microdrive presented in Section 4 (corresponding figures are respectively labeled [NAN drive] and [CIT drive]). We first provide a detailed look at the "isolation" process, and then consider algorithm performance over a broader time scale.

5.1. Neuron Isolation Process

Figure 8(a) summarizes algorithm behavior and recording quality during transitions through the prototypical SFSM path from *spike search* through *neuron isolated*, as the SNR curve is adaptively estimated and maximized. This plot displays electrode depth, the IQM and SNR metrics of the dominant neuron, and the state of the SFSM for 12 consecutive algorithm cycles, where each cycle consists of 10 seconds of recording, followed by about 25 seconds for analysis and electrode movement. Also, the recorded spike waveforms and their two-dimensional PCA projections are shown for selected time steps. The corresponding SNR curve is presented in Figure 8(b), where the horizontal axis represents electrode depth u.

After spikes are found at step $k = 2$, a positive gradient is found at step $k = 4$, and then the SNR curve is constructed, with the maximum reached by $k = 10$. The SFSM waits one cycle for confirmation before transitioning to *neuron isolated* at $k = 11$. Note that the spike waveform amplitudes are clearly improved by the electrode movement during steps $k = 2$ to $k = 11$. Also, the algorithm has tracked neuron "A" in PCA space. The algorithm continues to optimize the dominant neuron's signal quality, even though a sparsely firing neuron with slightly higher SNR is later identified (as seen in the detail of $k = 11$), since this is the neuron to which the SNR curve belongs and it is providing a strong signal.

Figure 9(a) shows data from a different representative experiment to illustrate an alternative transition to *neuron isolated*, as the signals of the dominant neuron invoke a jump in

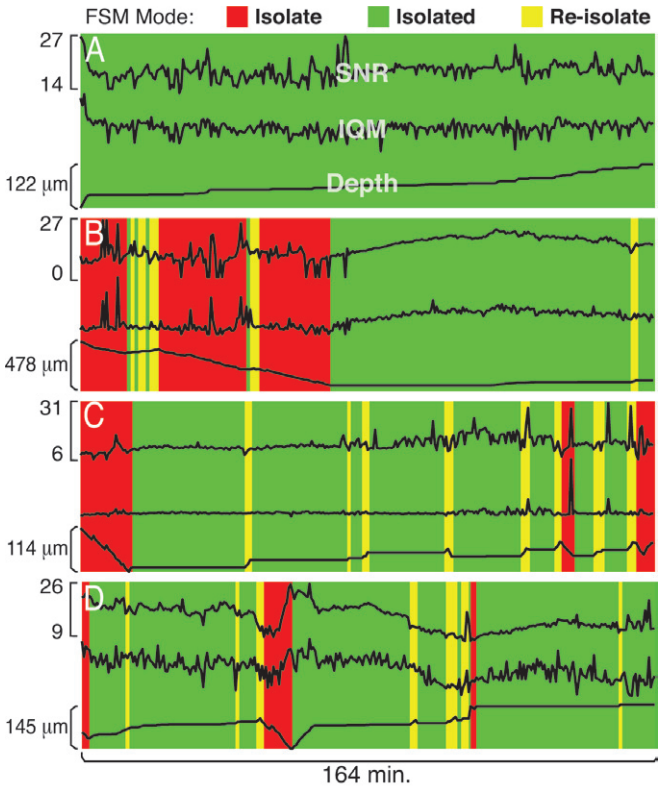

Fig. 10. SFSM mode, SNR, IQM, and electrode depth over time on four simultaneously operated electrodes during an acute recording session [NAN drive].

the IQM to interval Ω_3 in Table 1 (at $k = 7$, confirmed at $k = 8$). By viewing the spike waveforms, it is clear a strong new neuronal signal appears at step $k = 7$. Note that this transition was not triggered during the maximization of a SNR curve. The transition to *neuron isolated* bypassed the *isolate neuron* state completely. This figure also presents an example where *gradient search* does not identify a positive SNR slope (from $k = 2$ to $k = 6$) and so does not transition to *isolate neuron*. Figure 9(b) shows why a re-isolation pathway is necessary in the SFSM. Here, a strong isolation has been achieved prior to $k = 1$, and, although the electrode is held stationary, by step $k = 6$ the dominant neuron's SNR has drastically decreased, presumably due to tissue relaxation. The algorithm detects this drop and retracts the electrode to attempt re-isolation. During the retraction process, the signal SNR reaches a level equivalent to the original isolation and thus the SFSM returns to *neuron isolated*, having retracted about 20 μm.

5.2. Maintaining Isolations over Time

In this section we consider algorithm performance over broader time scales. For convenience of display and discussion, the states of the SFSM are aggregated into three *modes*: **isolate**, which includes *spike search*, *gradient search*, and *isolate neuron*; **isolated**, equivalent to *neuron isolated*; and **re-isolate**, consisting of *re-estimate gradient* and *re-isolate neuron*.

Table 2. Performance Metrics for One Month of Algorithm Use.

Electrode-hours under autonomous control	153
Percentage of time with neuron isolated	56%
Percentage of time attempting to isolate	32%
Percentage of time attempting to re-isolate	12%
Number of isolations per electrode per day lasting at least 30 minutes	1.2
Number of isolations per electrode per day lasting at least 60 minutes	0.65

Figure 10 provides the SFSM mode, SNR, IQM, and electrode depth for four simultaneously operated electrodes, representing scientific recording sessions from a particular day. Here, electrode A represents a nearly ideal case wherein a neuron is deemed to be isolated throughout the recording. The other electrodes show more typical cases: the algorithm first seeks neuronal signals, builds the SNR curve, isolates a neuron, and then intermittently readjusts electrode position to re-isolate the neuron. While electrode B does not find a stable isolation until about 80 minutes into the session, advancing nearly 500 μm before achieving, it also holds an isolation for the majority of the recording session.

Figure 11 summarizes the evolution of the algorithm's operating modes during 16 consecutive recording sessions, representing all algorithms used in one particular calendar month by a neurophysiologist. Each of the 16 sets of three or four bars represents all electrodes used on that day, in a manner similar to Figure 10, but without the line plots. Note that the recording sessions lasted for varying lengths of time, and missing bars indicate that the electrode was not used that day[7]. Some simple performance metrics for this data set are provided in Table 2. Of the trials that spend 75% or more time in "attempting to isolate", the time in this mode is broken down into: 64% in *spike search*; 27% in *gradient search*; and 8% in *isolate neuron*. This indicates that these sessions are unproductive primarily due to a lack of activity along the electrode's path (no detected spikes), as the electrode was placed in electrically inactive, or weakly active, neural tissue.

7. Days with missing bars typically indicate a broken electrode. Bars of varying length correspond to recording sessions where: (a) the electrode exhausted its range (and thus it was stopped and withdrawn) without finding a useful signal; (b) an electrode wire broke; or (c) the scientific goals of that recording were met, and the recording was discontinued by the neurophysiologist.

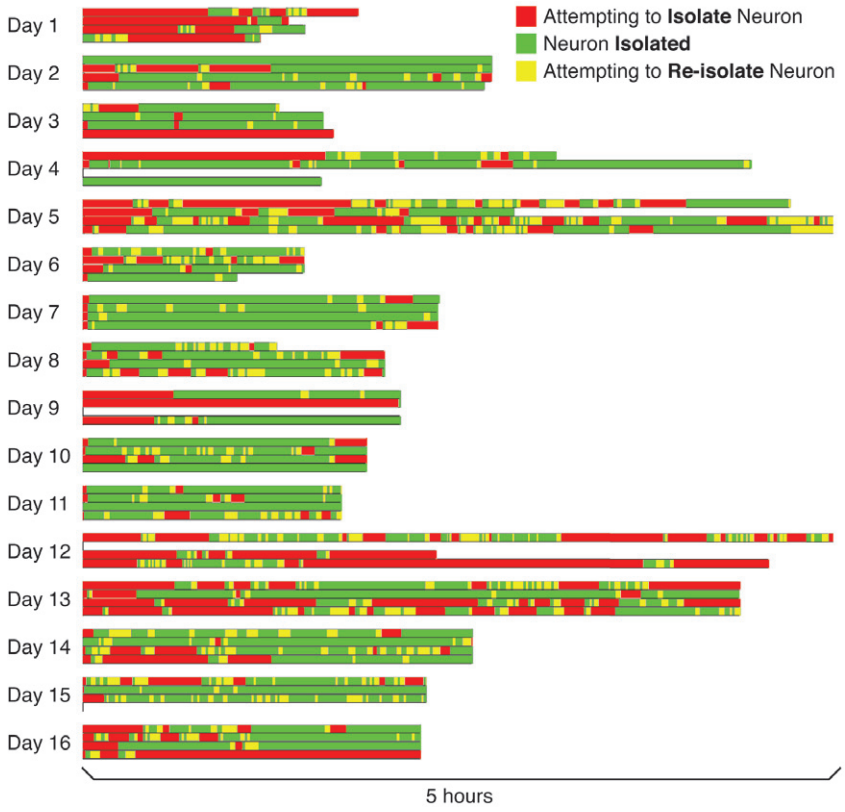

Fig. 11. SFSM mode for 16 consecutive recording sessions from one calendar month of algorithm use, typically operating four electrodes per day [NAN drive]. Length of session varies by day.

6. Conclusion and Future Work

In this paper we have described a paradigm for autonomous control of neural recording electrodes. The core of the control system consists of several algorithms within a hierarchical structure to process neuronal signals, optimize signal quality, and address practical recording challenges. This control system has been used to automate several different existing microdrives. We also presented a novel microdrive design intended as a mesoscale testbed for this control system. Experiments in primate cortex have validated our approach in acute settings. Beyond the results presented in Section 5.2, we note anecdotally that our control system has been used to autonomously control electrode movement for over 1,000 recording hours, achieving and maintaining hundreds of cell isolations in four different experimental facilities and multiple animal subjects. As human performance in neuron isolation has not been well characterized, a comparison between the algorithm and a human operator is not readily achieved, and would be complicated by the variability in recording quality across electrodes, areas of electrode insertion, recording hardware, and other factors.

We envision two primary future applications of this work. First, our results to date support the use of this technology to automate many aspects of the acute cortical recordings that are routinely carried out by neurophysiologists. In practice, a single neurophysiologist can at best manually control three to five electrode positions simultaneously during an acute recording experiment. However, multi-electrode recordings involving dozens of electrodes are of increasing scientific interest. By automating many of the manually laborious neural recording processes our algorithms could enable a single researcher to simultaneously manage a significantly larger number of electrodes during an acute recording, while still maintaining acceptable recording quality. In this way, our algorithm can potentially increase the productivity of acute neurophysiological experiments.

Second, we believe that our algorithms might enable improvements in future implantable neural interfaces used for neuroprostheses. The ability to continually and independently reposition the electrodes of an implanted array would improve the signal quality and yield of signals as compared with conventional fixed geometry implanted arrays. Preliminary experiments also suggest that the ability to continually reposition an

electrode increases the neuroprosthetic "value" of the recorded data, as compared with current fixed geometry arrays. In these experiments, the "decoding efficiency" (roughly, the contribution of each electrode's signal to the task of decoding the neuroprosthetic user's intent) obtained with continually repositioned electrodes increased, relative to that of a fixed electrode array, by a factor of two to three (Mulliken 2008). These results imply that future chronic arrays with autonomous electrodes may increase not only the number of electrodes recording active neurons but also perhaps the "value" of the data obtained by each electrode. Future improvements in the control algorithm could even further increase this value by specifically seeking out neurons whose activity is strongly correlated with the objectives of the neuroprosthesis. However, further studies (which are feasible with the prototype presented in Section 4) on the longer-term effects of electrodes' frequent motion in cortical tissue are needed. In ongoing work we are developing new microelectromechanical system (MEMS) actuators and fabrication methods to build implantable arrays of many individually actuated recording micro-electrodes (Pang et al. 2006).

Appendix: SFSM Mechanisms to Cope with Additional Neural Recording Difficulties

Additional issues that typically arise in practical neural recording situations, and the SFSM mechanisms that cope with each issue, are described in this appendix.

Non-stationarity Owing to Tissue Decompression

At the start of an acute experiment, before the fragile electrodes are advanced into the brain, protective "guide tubes" (see Section 4) must puncture through the dura (the tough protective layer between the brain and skull). This process causes compression of brain tissue, and further compression may occur as the electrodes travel to their desired depth. Once the electrodes stop moving, the tissue decompresses over a several-hour period. Smaller scale tissue drift can also occur following optimizing electrode movements: evidence of such drift is observed after electrode movements of less than 100 μm. The SFSM accounts for this non-stationarity via the "re-isolate" states and, in more drastic cases, by initiating a search for a new neuron if necessary. Decompression effects also factor into some of the issues described in the following.

Cell Death

The moving electrode can potentially impale a neuronal cell body, causing its death. This possibility motivates the definition of IQM interval Ω_3 in Table 1. Alternatively, the electrode must retract as tissue decompression may carry a cell body towards a stationary electrode. As excessive SNR is a good indication of this situation, the electrode is retracted when the SNR value exceeds a maximum threshold, $\text{SNR}_{\max}$. The amount of retraction is proportional to the extent to which the SNR exceeds $\text{SNR}_{\max}$, thereby effecting a proportional control loop. This BACK AWAY SFSM state may be invoked at any time.

Discrete/Transient Events

The SFSM is designed to limit sensitivity to transient events that may arise from a number of sources.

- **Intermittent Neural Activity.** Individual neurons may lapse into brief periods of inactivity during which no spikes are emitted. When the algorithm attempts to isolate (or has isolated) a neuron that becomes silent over T_k, it may attempt to isolate a different neuron instead. Often, the neuron will resume firing on subsequent intervals.
- **False Spikes.** A variety of spurious noise artifacts may lead to recorded signals that appear quite similar to spikes. For example, abrupt movement of the subject can cause the electrode to vibrate in its guide tube, yielding transient artifacts. Electromagnetic fields emitted from the electrode drive motors can also introduce spike-like waveforms. These events may be falsely detected by the spike detection module and clustered separately as a new dominant target neuron, or clustered with the spikes of the existing dominant neuron (grossly affecting the signal metrics for that sample).
- **Uncommanded Electrode Shifts.** Large subject movements can induce sudden electrode displacement relative to the neural tissue, causing a well-isolated neuron signal to disappear, or making it impossible to track neurons across recording intervals.

Experience has shown that these and similar cases can be handled via the use of a WAIT state (see Figure 4) that delays a transition until the transition criteria is met for R consecutive recording intervals (where $R = 2$ in our experiments). For example, if a neuron were to cease spiking while the SFSM is in *isolate neuron*, the SFSM will effectively hold the SNR curve for $(R - 1)$ additional cycles before transitioning back to *spike search*.

Electrode–Tissue Mechanical Interactions

There is often a lag between electrode movements and the ensuing changes in signal quality. Such hysteresis may arise

from friction between the electrode tip and the surrounding tissue, causing a stick–slip delay between an electrode movement command and the actual displacement with respect to nearby cell bodies. Hysteresis may also arise when the electrode changes direction. It may be caused by the same frictional effect, or the cell bodies of nearby neurons may be shifted by the tissue displacement caused by the electrode's movement. For these and other sources of hysteresis and nonstationarity, the signal quality is monitored as a function of time as well as electrode depth. A consistent downward trend in signal quality, regardless of the direction of electrode travel, signifies that the signal quality estimates cannot be trusted. In these cases, if the dominant neuron is acceptably isolated (IQM $\in \Omega_2$), then a transition to *neuron isolated* is invoked, as the data trend suggests that further electrode movements will not improve signal quality. If the signal quality trends downward and the isolation is *not* acceptable (IQM $\in \Omega_1$), the SFSM triggers *spike search* in order to find a different neuron.

References

Andersen, R. A., Burdick, J. W., Musallam, S., Pesaran, B., and Cham, J. G. (2004). Cognitive neural prosthetics. *Trends in Cognitive Sciences*, **8**(11): 486–493.

Avezaat, C. J. J. and van Eijndhoven, J. H. M. (1986). The role of the pulsatile pressure variations in intracranial pressure monitoring. *Neurosurgical Review*, **9**(1): 113–120.

Baker, S. N., Philbin, N., Spinks, R., Pinches, E. M., Wolpert, D. M., MacManus, D. G., Pauluis, Q. and Lemon, R. N. (1999). Multiple single unit recording in the cortex of monkeys using independently moveable microelectrodes. *Journal of Neuroscience Methods*, **94**(1): 5–17.

Branchaud, E. A. (2006). An algorithm for the autonomous isolation of neurons in extracellular recordings. *PhD Thesis*, California Institute of Technology.

Branchaud, E. A., Andersen, R. A. and Burdick, J. W. (2006). An algorithm for autonomous isolation of neurons in extracellular recordings. *IEEE/RAS-EMBS International Conference on Biomedical Robotics and Biomechatronics (BioRob)*.

Branchaud, E. A., Cham, J. G., Nenadic, Z. and Burdick, J. W. (2005). A miniature robot for autonomous single neuron recordings. *IEEE International Conference on Robotics and Automation (ICRA)*.

Buzsaki, G. (2004). Large-scale recording of neuronal ensembles. *Nature Neuroscience*, **7**(5): 446–451.

Cham, J. G., Branchaud, E. A., Nenadic, Z., Greger, B., Andersen, R. A. and Burdick, J. W. (2005). Semi-chronic motorized microdrive and control algorithm for autonomously isolating and maintaining optimal extracellular action potentials. *Journal of Neurophysiology*, **93**: 570–579.

Cham, J. G., Wolf, M. T., Andersen, R. A. and Burdick, J. W. (2006). Miniature neural interface microdrive using parylene-coated layered manufacturing. *IEEE/RAS-EMBS International Conference on Biomedical Robotics and Biomechatronics (BioRob)*.

deCharms, R. C., Blake, D. T. and Merzenich, M. M. (1999). A multielectrode implant device for the cerebral cortex. *Journal of Neuroscience Methods*, **93**(1): 27–35.

Emondi, A. A., Rebrik, S. P., Kurgansky, A. V. and Miller, K. D. (2004). Tracking neurons recorded from tetrodes across time. *Journal of Neuroscience Methods*, **135**: 95–105.

Fee, M. S. (2000). Active stabilization of electrodes for intracellular recording in awake behaving animals. *Neuron*, **27**(3): 461–468.

Fee, M. S. and Leonardo, A. (2001). Miniature motorized microdrive and commutator system for chronic neural recording in small animals. *Journal of Neuroscience Methods*, **112**(2): 83–94.

Fraley, C. and Raftery, A. E. (1998). How many clusters? Which clustering method? Answers via model-based cluster analysis. *Computer Journal*, **41**(8): 578–588.

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

Harris, K. D., Hirase, H., Leinekugel, X., Henze, D. A. and Buzsaki, G. (2001). Temporal interaction between single spikes and complex spike bursts in hippocampal pyramidal cells. *Neuron*, **32**(1): 141–149.

Kralik, J. D., Dimitrov, D. F., Krupa, D. J., Katz, D. B., Cohen, D. and Nicolelis, M. A. L. (2001). Techniques for long-term multisite neuronal ensemble recordings in behaving animals. *Methods*, **25**(2): 121–150.

Lewicki, M. S. (1998). A review of methods for spike sorting: the detection and classification of neural action potentials. *Network: Computation in Neural Systems*, **9**: R53–R78.

Mulliken, G. H. (2008). Continuous sensorimotor control mechanisms in posterior parietal cortex: forward model encoding and trajectory decoding. *PhD Thesis*, California Institute of Technology.

Muthuswamy, J., Okandan, M., Gilletti, A., Baker, M. S. and Jain, T. (2005). An array of microactuated microelectrodes for monitoring single-neuronal activity in rodents. *IEEE Transactions on Biomedical Engineering*, **52**(8): 1470–1477.

Nenadic, Z. and Burdick, J. W. (2005). Spike detection using the continuous wavelet transform. *IEEE Transactions on Biomedical Engineering*, **52**(1): 74–87.

Nenadic, Z. and Burdick, J. W. (2006). A control algorithm for autonomous optimization of extracellular recordings. *IEEE Transactions on Biomedical Engineering*, **53**(5): 941–955.

Pang, C., Tai, Y. C., Burdick, J. W. and Andersen, R. A. (2006). Electrolysis-based diaphragm actuators. *Nanotechnology*, **17**(4): S64–S68.

Porada, I., Bondar, I., Spatz, W. B. and Kruger, J. (2000). Rabbit and monkey visual cortex: more than a year of recording with up to 64 microelectrodes. *Journal of Neuroscience Methods*, **95**(1): 13–28.

Rousche, P. J. and Normann, R. A. (1998). Chronic recording capability of the Utah intracortical electrode array in cat sensory cortex. *Journal of Neuroscience Methods*, **82**(1): 1–15.

Serruya, M. D., Hatsopoulos, N. G., Paninski, L., Fellows, M. R. and Donoghue, J. P. (2002). Brain–machine interface: instant neural control of a movement signal. *Nature*, **416**(6877): 141–142.

Taylor, D. M., Tillery, S. I. H. and Schwartz, A. B. (2002). Direct cortical control of 3D neuroprosthetic devices. *Science*, **296**(5574): 1829–1832.

Turner, J. N., Shain, W., Szarowski, D. H., Andersen, M., Martins, S., Isaacson, M. and Craighead, H. (1999). Cerebral astrocyte response to micromachined silicon implants. *Experimental Neurology*, **156**(1): 33–49.

Venkatachalam, S., Fee, M. S. and Kleinfeld, D. (1999). Ultraminiature headstage with 6-channel drive and vacuum-assisted micro-wire implantation for chronic recording from the neocortex. *Journal of Neuroscience Methods*, **90**(1): 37–46.

Vos, B. P., Wijnants, M., Taeymans, S. and De Schutter, E. (1999). Miniature carrier with six independently moveable electrodes for recording of multiple single-units in the cerebellar cortex of awake rats. *Journal of Neuroscience Methods*, **94**(1): 19–26.

Wessberg, J., Stambaugh, C. R., Kralik, J. D., Beck, P. D., Laubach, M., Chapin, J. K., Kim, J., Biggs, S. J., Srinivasan, M. A. and Nicolelis, M. A. L. (2000). Real-time prediction of hand trajectory by ensembles of cortical neurons in primates. *Nature*, **408**(6810): 361–365.

Williams, J. C., Rennaker, R. L. and Kipke, D. R. (1999). Long-term neural recording characteristics of wire microelectrode arrays implanted in cerebral cortex. *Brain Research Protocols*, **4**(3): 303–313.

Wolf, M. T. and Burdick, J. W. (2007). Spike clustering and neuron tracking over successive time windows. *IEEE EMBS Conference on Neural Engineering (NER)*.