On improving the effectiveness of control signals from chronic microelectrodes for cortical neuroprostheses

by

Hirak Parikh

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Biomedical Engineering)
in The University of Michigan
2009

Doctoral Committee:

Professor Daryl R. Kipke, Chair
Assistant Professor Joshua Damien Berke
Assistant Professor Edward L. Ionides
Assistant Professor Parag G. Patil

Dedicated to my parents:

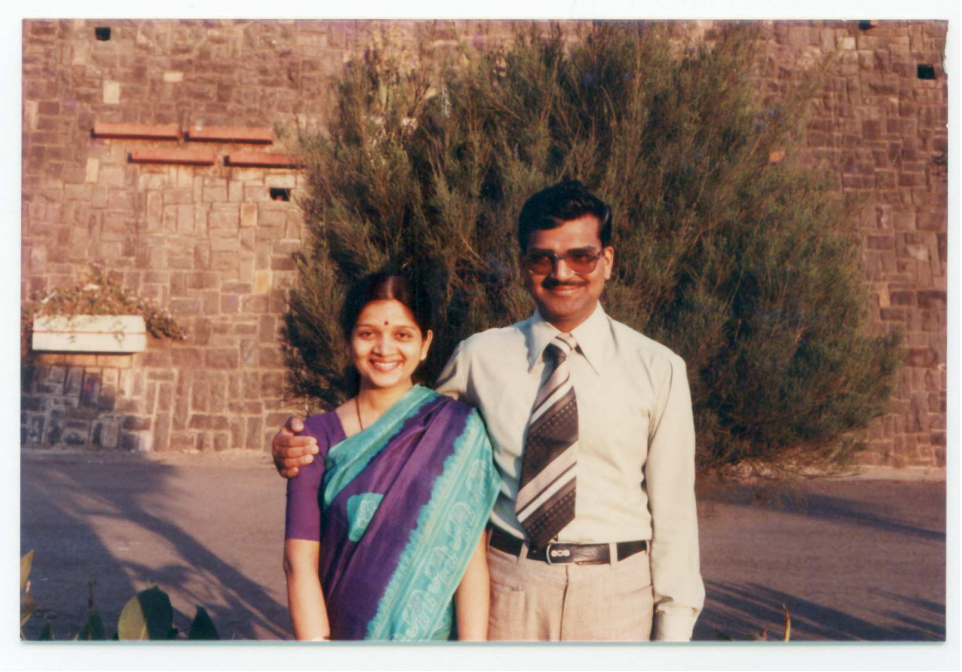

Dr. Sunit and Pragnya Parikh,
whom I can never thank enough.
This dissertation is submitted in partial fulfillment of my debt to you.
I hope, I continue to do you proud.

ACKNOWLEDGEMENTS

Daryl Kipke, my mentor and advisor who let me in his laboratory and then let me loose to pursue my projects. Thanks for guiding me and even prodding me when the occasion arose.

Joshua Berke, who always raised the bar. Thanks! it was all worth it in the end.

Tim Marzullo, who has been a fantastic colleague, co-author, co-adventurer, and co-conspirator through all these years in grad school. It's been great to share joys and disappointments, hobbies and your home-made brew.

Greg Gage, another co-conspirator and fellow caffeine addict whose figures and code will always be better than everyone else's. Thanks for your manuscript comments and funny stories.

Mark Lehmkuhle, who diligently read every draft of all my manuscripts and never tired of telling me to increase the font size in the figures. Thanks for your espresso machine whose regular deliverance of caffeine fueled much of the following work.

Also, thanks to all the other NEL rats: Buckeye John, Kip, Jey, Nick, Matt J., Matt G., Rachel, Erin, Eugene, Azadeh, Mohammed, Ninja T.K., Pratik and Taegyun for helping me at various stages, including when I got locked out.

Special thanks to recent collaborators – Anindya Bhadra, Alex Wiltschko and Chun-hui Chang.

Thanks also to Tonya Brown, Vera Williams and Karen Coulter in the office.

The City of Ann Arbor, which has been my first home in this land of the free and

home of the brave. Thanks for the beautiful running trails, parks, summer concerts, art fairs, bars, and restaurants.

My friends from the Quiz Club and Michigan Argentine Tango Club, who have provided entertainment and escape. You helped me retain my sanity during these last few years.

My extended family scattered across the globe, who are wondering what I've been up to lately (last half a dozen years, or so). Thanks for your support. Now, I hope it all makes sense; see attached explanation.

Sumedha Sinha, who has been my most exacting editor, fiercest critic and steadiest supporter. Who has endured my absence both mental and physical in the last year. Thanks not only for being my wife and best friend, but also for just being there.

TABLE OF CONTENTS

LIST OF FIGURES

Figure

LIST OF TABLES

Table

ABSTRACT

On improving the effectiveness of control signals from chronic microelectrodes for cortical neuroprostheses

by
Hirak Parikh

Prof. Daryl R. Kipke

Using microelectrodes, we can record neural signals which can eventually be used to control cortical neuroprostheses for assisting people with spinal-cord trauma, stroke deficits, amyotrophic lateral sclerosis (ALS), and motor-neuron disease. Despite recent encouraging advances, a number of fundamental issues need to be resolved for a reliable, fully-functional, long-term human neuroprosthesis. Improved cortical prostheses require further development both in neural interfaces and investigation of cortical signals for obtaining the most effective control signals. The goal of this dissertation is to investigate the effectiveness of unit activity and local field potentials (LFPs) in the motor cortex using chronic multisite microelectrodes.

In the first study, we first demonstrate a novel method to assess neural signatures across sessions and quantify neuron stability by providing a probabilistic estimate of similarity between spike clusters. This technique supports both single and multiple electrodes, and has applications in designing appropriate neuroprosthetic control algorithms, determining recalibration parameters, investigating neural plasticity, and

assessing significance of particular metrics.

Next, we investigate unit activity and LFP activity in the different layers of the motor cortex. Four rats were implanted bilaterally with multi-site single-shank silicon microelectrode arrays in the motor cortex while the animal was engaged in a movement-direction task. In the second study, we demonstrate that units in the lower layers (Layers 5,6) are more likely to encode direction information as compared to units in the upper layers (Layers 2,3) suggesting electrode sites clustered in the lower layers provide access to more salient control information.

In the third study, we investigate LFP activity to determine significant interactions in time and/or frequency across the different layers. We analyzed LFP activity in four frequency ranges: low (3-15Hz), low-gamma (15-40Hz), high-gamma (40-70Hz) and high (>70Hz) across both upper (Layers 2,3) and lower layers (Layers 5,6) of the cortex. Our analysis based on 585 LFP recordings from 39 sessions shows that the low frequency range (3-15Hz) is more likely to encode directional information as compared to other frequency ranges. We found a significant difference in LFP activity between the upper and lower layers of cortex in the high gamma (40-70Hz) range, but not in the other frequency ranges. Our results indicate that LFPs are viable alternative control signals that can be recorded from either upper or lower layers of the cortex for performance comparable to our results from unit activity.

CHAPTER I

Introduction

The brain is the fount of all reason. It is the initiator and mediator of movements, feelings, and perceptions. It is the most complex organ of the human body and also the least understood. Unlike other cells in the body, in general, neurons do not regenerate. Injury to brain does not always cause long-term impairment or disability. Depending on the location and extent of damage, there can be personality changes, neurocognitive deficits, speech, movement, and mental handicaps. Severe brain damage may result in persistent vegetative state, coma, or death.

The use of electrical stimulation has been used to influence brain function and treat neurological disorders since the 1950s [Cooper, 1981],[Delgado, 1967],[Delgado, 1969]. Physical devices have been implanted in the brain to treat neurological disorders such as Parkinson's disease, spinal cord injuries, degenerative muscular diseases, stroke or other nervous system injury [McLachlan, 1997],[Muir and Steeves, 1997],[Benabid et al., 2001]. Using microelectrode technology, we can record neural signals which can be used to control cortical neuroprostheses, or brain-computer interfaces (BCIs), by processing the recorded neural signal and extracting a control signal which then operates an external device [Schwartz, 2004]. Indirect or non-invasive neural control of a cursor

by a human has been demonstrated by [Wolpaw and McFarland, 2004]. Advances have been made in building such neuroprosthetic systems in the motor cortex [Nicolelis, 2001], [Serruya et al., 2002], [Taylor et al., 2002], auditory [Loeb, 1990], and visual cortex [Maynard et al., 1999]. Successful 2-D cursor control was demonstrated in humans in 2006 [Hochberg et al., 2006], showing the promise of these technologies for people paralyzed by spinal-cord trauma, stroke deficits, amyotrophic lateral sclerosis (ALS), cerebral palsy, multiple sclerosis, and motor-neuron disease.

These recent advances in techniques and methods are indeed very encouraging, but a number of fundamental issues need to resolved before we have a reliable, fully-functional, long-term human neuroprosthesis. Direct cortically-derived command signals have shown encouraging results with non-human primates, but the same level of control has not been achieved in humans and current systems are much slower as compared to natural human behavior, with limited dimensionality [Donoghue, 2002]. The basic requirement for neuroprosthetic devices is the ability to record unit activity reliably for many years and the ability to record action potentials from many different cells in a small volume of cortex [Schwartz, 2004]. There are limitations at the level of the neural interface: the amount of tissue damage that is acceptable and the longevity of the device and the signals that can be recorded from the devices. There have been advances in reliable chronic multi-electrodes and interfaces, but there still remains a lack of suitable techniques to analyze the signals from the brain and model its function effectively to obtain the optimal control signals for a human neuroprosthesis.

1.1 Overview

The chief goal of this dissertation is to investigate the effectiveness of cortical signals using chronic microelectrodes in the motor cortex in a rat model. We first demonstrate a method to assess and track neurons across multiple sessions to better characterize their activity and quantify their stability. Second, we investigate two kinds of cortical signals – firing of individual neurons (units) and local field potentials (LFPs) – in the different layers of the motor cortex and determine the most effective modality/parameters for control signals. What follows in this chapter is the scientific and technological background, followed by an introduction and experimental justification for the three studies comprising this dissertation.

1.2 How? Recording brain signals

Brain activity can measured by recording electric potentials caused by the activity of individual neurons at different spatial and frequency resolutions. Intracellular recordings allow monitoring of activity of a single neuron; but are not feasible for long-term recordings in awake and behaving preparations. Extracellular recording technologies allow us to chronically record from multiple cells in awake and behaving preparations by placing an electrode close to the cell body. In addition to recording the activity of neurons, we can also record extracellular continuous signals such as local field potentials (LFPs), electrocorticograms (ECoGs), and electroencephalograms (EEGs).

For recording electrodes there are three main design considerations: 1) adequate spatial and temporal resolution of desired signals, 2) number, location and placement of the electrode sites, and 3) functional lifetime of the device [Kipke et al., 2008]. Different electrode technologies have been devel-

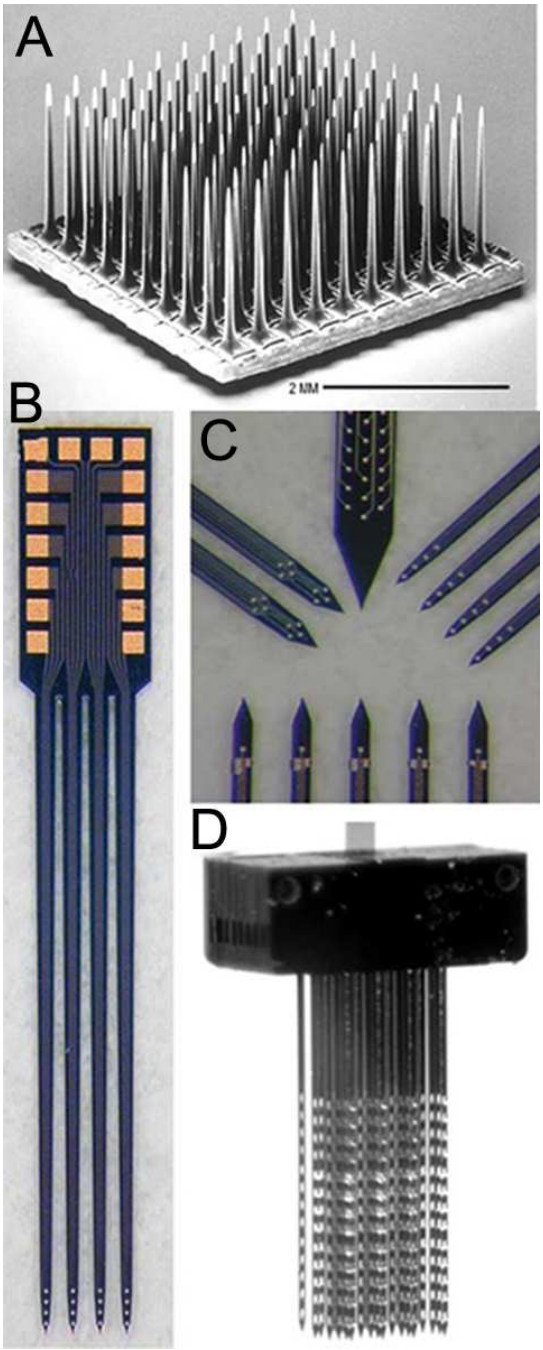

Figure 1.1: Microelectrode technologies: A) Utah probe showing a 100-electrode array with each electrode separated by 400 μm [Donoghue, 2002]. B) Michigan four-shank probe with four electrode sites at the tip C) High-magnification photographs illustrating four different types of sites layouts for specialized interfaces D) Modular 128-site, three dimensional array [Kipke et al., 2008].

oped to meet these needs [McNaughton et al., 1983] [Nordhausen et al., 1996], [Hetke JF, 2002], [Nicolelis et al., 2003], [Wise et al., 2004], [Rennaker et al., 2005], [McCreery et al., 2006], [Neves and Ruther, 2007] which enable simultaneous recording of multiple signals using multi-site electrodes at multiple spatial scales ranging from synaptic currents to single units to local fields.

One of the oldest techniques is using insulated microwires to record brain activity. Microwires are the least expensive and relatively easy to build, but cannot be manufactured with reliable geometries and their relatively large size leads to more tissue damage and deterioration in recording capability [Williams et al., 1999]. One kind of silicon microelectrode, popularly called the 'Utah probe', consists of a grid of finely spaced microneedles as shown in Figure 1.1 A. These probes have shown efficacy for long-term use [Schwartz, 2004], with consistency of the electrode manufacture with flexible geometries. The most flexible of these silicon-based technologies is the 'Michigan probe' which uses lithographic patterning of thin films of conductors and insulators on silicon or polymer substrates. As shown in Figure 1.1 B, these probes are highly flexible in terms of electrode site configuration with single or multiple shanks and can be fabricated from a variety of materials depending on the location of the implant and type of signal [Hetke JF, 2002]. Given the flexibility in design geometry it is possible to use an open architecture to minimize tissue damage [Seymour and Kipke, 2007]. Massively parallel recordings have been demonstrated using using these silicon electrodes [Csicsvari et al., 2003] and the longevity of these probes has been shown to be on the scale of months to a year [Vetter et al., 2004]. One of the chief advantages of the multi-site 'Michigan probes' is that we can simultaneously record brain activity at different cortical depths in awake and behaving animals in chronic preparations.

1.3 What? Different kinds of cortical signals

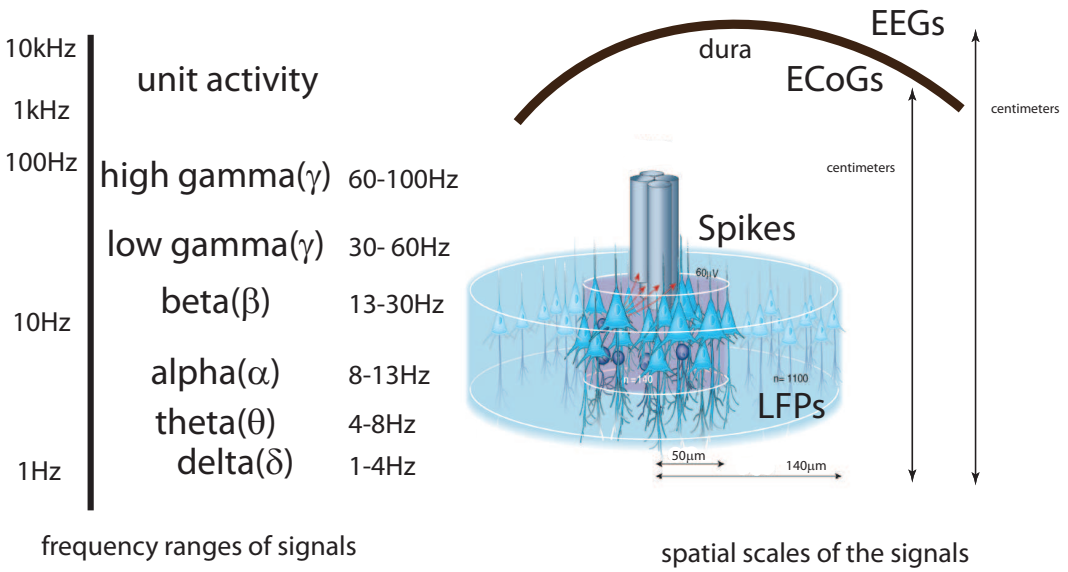

Figure 1.2: Brain Signals: The signals that can be recorded depend on the frequency spectrum being analyzed. Spikes are discrete signals that can be obtained by high pass filtering the recorded signal from 1kHz to 10kHz. Continuous signals such as LFPs, EEGs and ECoGs are much lower in frequency. They can be further split into different frequency sub-bands such as α(7-10Hz), β(18-24Hz) and γ(40-100Hz) waves. Ranges as per [Asher et al., 2007]. Right panel shows the spatial scales of the signals. The recording sphere of an electrode picks up spiking activity at a distance of tens of microns. LFPs are summed potentials averaged over 140μm. ECoGs are recorded subdurally by placing electrodes over the surface of the brain and EEGs are recorded from the scalp. Both signals represent average neural activity over centimeters. Adapted from [Buzsaki, 2004]

The brain is a rich source of both analog (field potentials, ECoGs, EEGs, etc.) and discrete signals, such as the firing of individual neurons. The type of signal obtained depends on the volume of brain tissue sampled, or appropriate separation of the frequency bands as shown in Figure 1.2.

Unit Activity: The firing of a neural action potential is an all-none event represented by a characteristic biphasic waveform (called a spike) as shown in Figure 1.3 A. Using microelectrode recordings we can record the activity of an individual neuron (also referred to as a 'unit'), or a small population of neurons. These spikes are obtained from the continuous recording by high-pass filtering the signal between

1kHz to 10kHz. Waveforms of one neuron are distinguished from the noise and other neurons by using a variety of techniques [Lewicki, 1998], [Harris et al., 2000]. The chief advantage of recording unit activity is that these recordings are highly specific and directly represent neuron activity. This technique is invasive, requiring the insertion of the microelectrodes into the cortex, and it results in tissue response and glial scarring. From the perspective of neuroprosthetics this leads to a difficulty in reliable long-term recordings for a number of years. Recording from the same cell over multiple sessions is problematic and would require recalibration of prosthetic control algorithms [Andersen et al., 2004].

Hence, recent studies have looked at alternative sources for obtaining control information for a neuroprosthetic device described below:

EEGs and ECOGs: Electroencephalograms (EEGs) and electrocorticograms (ECoGs) are brain waves recorded from the scalp and brain surface respectively and result from mass changes in population synaptic activity from the cerebral cortex. The advantage of the EEG signal is that it is robust over time and is recorded non-invasively. EEG recordings can detect events lasting only a few thousandths of a second which make them attractive for human neuroprostheses. Compared to penetrating electrodes, ECoGs are less invasive since they record from the surface of the brain. The relationship between movements and fast oscillations of the ECoG and EEG has been investigated in humans [Aoki et al., 1999], [Pfurtscheller et al., 2003], [Leuthardt et al., 2004]. It has been shown that EEG and ECoG activity can control a simple brain-machine interface [Wolpaw, 2004]. ECoG-based BCIs have been shown to have a success rate of 55-73% in a two-dimensional, four-target center-out task [Leuthardt et al., 2004], [Schalk et al., 2007],[Schalk et al., 2008]. The chief disadvantage is that EEGs are comprised of signals summed over centimeters of brain,

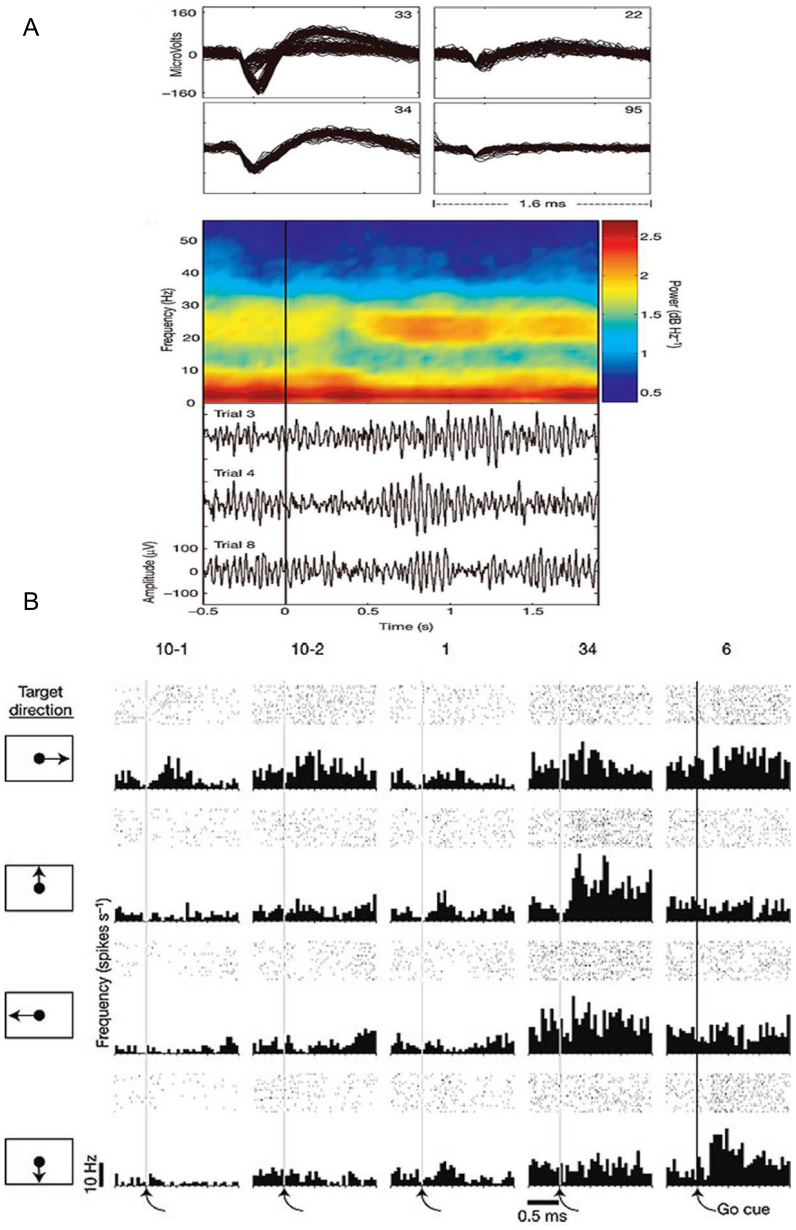

Figure 1.3: Examples of unit activity and local field potentials (LFPs) obtained from the first human implant. A) Top panel shows typical biphasic waveforms of neurons obtained from four different electrode sites. Middle panel shows average local field potential (LFP) activity for one sample signal shown as a function of time and frequency (spectrogram). Bottom panel shows evoked LFP activity in three different trials in the time domain. B) Raster plots and peri-stimulus time histograms showing unit activity in response to 'Go' cue of five different neurons recorded in first human neuroprosthesis. As the subject imagined movements to four different directions, brain activity was decoded to control the movement of a cursor on the screen. [Hochberg et al., 2006].

and thus have limited specificity, with very little sensitivity for recording sub-cortical activity.

LFPs: The EEG and single-cell recordings sum activity over areas of very different scale: over centimeters for the EEG and over microns for cell recording. The LFP lies between these two scales of sampled activity. This signal comprises the activity of hundreds or thousands of cells around an electrode tip inserted into the cortex or placed on the cortical surface. LFPs are assumed to be largely a result of summed electric potentials resulting from excitatory post-synaptic potentials (EPSPs) and synaptic activity [Mitzdorf, 1985]. Like single-cell recordings, LFP recordings are invasive. The 'listening sphere' of LFPs is large as they are less affected by scarring and they can be recorded for a longer time as compared to unit activity [Buzsaki, 2004]. A recent study estimated the recording sphere of LFPs to be within 250μm of the recording microelectrode [Katzner et al., 2009]. Recorded spiking activity is biased towards the activity of larger cells, which are more likely to have connections with other brain areas, whereas LFPs are generated by local synaptic activity. LFP oscillations have coincided with epochs of increased or decreased neural discharges around movement in the motor cortex [Donoghue et al., 1998], and have also been shown to encode for behavioral states and movement direction [Mehring et al., 2003, Rickert et al., 2005, Scherberger et al., 2005, O'Leary and Hatsopoulos, 2006]. Thus, LFPs can be used as additional sources of control information to augment the usable lifetimes of microelectrode implants as they are more stable than unit activity [Andersen et al., 2004].

It was shown that by combining the activity of single units and LFPs better decoding performance was obtained than using the signals alone [Mehring et al., 2003]. Hence, to achieve long-term recordings and for greater degrees of freedom, future

systems will need to rely on a number of different brain signals and maximize the theoretical control information encoded in these signals.

1.4 Where? Locations for a neuroprosthesis

The primary target for a neuroprosthetic device has been the motor cortex, since a collection of motor cortical areas contribute to performing voluntary movement [Donoghue and Wise, 1982], [Murthy and Fetz, 1996], [Donoghue et al., 1998]. In case of a severed spinal cord or injury to motor neurons, the motor cortex is capable of still generating signals [Shoham et al., 2001]. A number of studies have verified that neural signals in the motor cortex encode control information that can be used to control a neuroprosthetic device [Georgopoulos et al., 1989] [Taylor et al., 2002], [Velliste et al., 2008]. In addition, a number of other brain areas have also shown promise for recording control signals such as the lateral inter-parietal area (LIP) [Andersen et al., 2004, Scherberger et al., 2005], cingulate [Marzullo et al., 2006], PmD [Achtman et al., 2007],[Hatsopoulos et al., 2004], etc.

Structurally, the neocortex has a regular vertically-oriented pattern consisting of six layers. While the relative thicknesses of the layers and number of neurons differ depending upon the area of the brain [DeFelipe et al., 2002], the overall structure is remarkably similar across brain areas. This suggested investigation into whether this structure had any functional significance [Mountcastle, 1997]. Observations in the cat somatosensory cortex led Mountcastle and others to hypothesize the concept of a minicolumn [Mountcastle, 2003]. According to this theory, each minicolumn extends perpendicularly across all six layers and forms a basic functional unit. These minicolumns form a subset of larger functional units called macrocolumns and columns. The basic concept has been verified by other

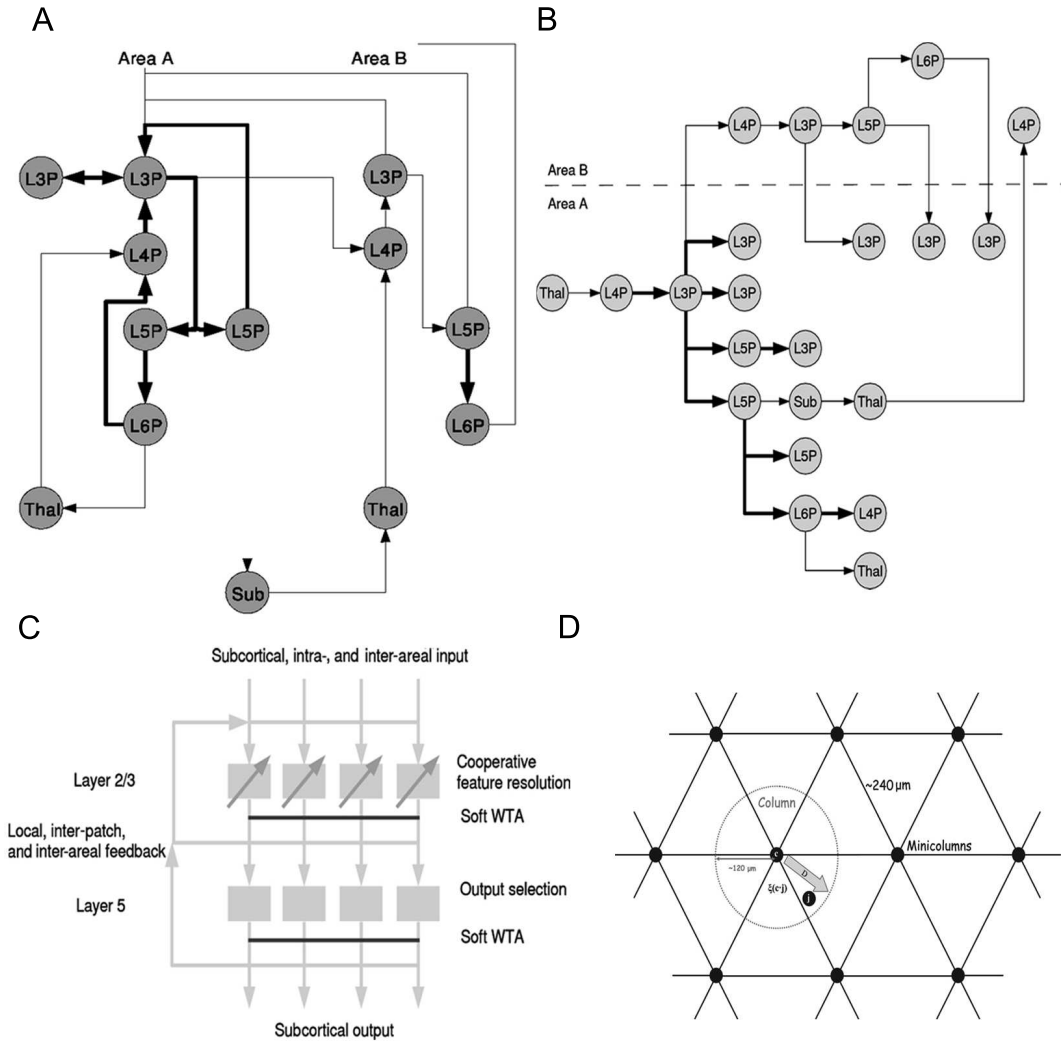

Figure 1.4: A) Hypothesized cortical microcircuits showing feedforward and feedback interactions between significant excitatory cells in the sub-cortical structures and the different layers of the cortex. Nodes are organized spatially; vertical corresponds to the layers of the cortex and horizontal to its lateral extent. B) Hypothesized temporal sequence of activity between cells across the different layers of the cortex and sub-cortical structures. Each edge represents one synaptic delay. C) A simple model of cortical processing that incorporates main aspects of cortical circuits. Neurons in the upper layer integrate sub-cortical, intra-areal and inter-areal input and participate in a selection network and cooperate to resolve a consistent interpretation; neurons in the lower layers process signals from the superficial signals and decide on the final output to motor structures [Douglas and Martin, 2004]. E) Cortical minicolumns are vertical structures that are perpendicular to the surface of the brain and span different layers of the cortex. In the motor cortex, cells in a particular minicolumn have a similar preferred direction; minicolumns in the motor cortex repeat at an approximate radial distance of 240μm [Georgopoulos et al., 2007].

researchers, however the size and thickness of minicolumns and macrocolumns have been debated [Rockland and Ichinohe, 2004]. One unresolved question is how this six-layered structure gives rise to function. A number of theoretical models have been posited on the connectivity of the cortex and how the different layers interact [Douglas and Martin, 2004]. Figure 1.4 A, B show the hypothesized connections between layers and flow of information from sub-cortical and thalamic regions to the various layers of the cortex respectively. According to this theory, neurons in the upper layers of the cortex integrate cortical input from sub-thalamic, intra-areal, inter-areal regions different areas and dynamically adjust their firing properties cooperatively select an interpretation or feature that is consistent with the different inputs. Based on activity from the superficial layers, neurons in the lower layers select the final output and drive the activity of sub-cortical structures [Douglas and Martin, 2004] as shown in Figure 1.4 C. The processing in the superficial layers is constrained via feedback connections from layer 5 pyramidal cells. Layer 5 pyramidal cells also project to the superficial layers of other cortical areas and provide additional context for cortical processing in other cortical areas. A recent study supported aspects of this hypothesis by showing that stimuli targeting the neurons upper layers evoked network-wide events, implying that upper layer neurons drive output neurons in lower layers [Weiler et al., 2008].

The functional significance of the layers has been more extensively studied in the sensory areas [Albright et al., 1984],[Bauer et al., 1983] than in the motor cortex. There have been recent efforts to investigate columnar and laminar organization in the motor cortex. It was found that M1 cells with similar preferred directions aggregate in a vertical dimension forming an ordered structure of minicolumns, perpendicular to the surface of the neocortex, of width $\approx 30\mu$m and repeating at a lateral

distance of 200μm [Amirikian and Georgopoulos, 2003], [Georgopoulos et al., 2007] as shown in Figure 1.4 D. A recent study found that dynamic local inhibition is responsible for altering the firing pattern of pyramidal cells that are narrowly tuned and present across all the layers of the motor cortex [Merchant et al., 2008]. They also discovered a second class of putative pyramidal neurons that are more broadly tuned and chiefly located in Layers 5 and 6.

1.5 How? Relationships between units and LFPs in time and frequency

It was shown that there was non-oscillatory synchrony among units during phasic movements when the requirement was to produce a highly fractionated pattern of muscle activity [Bennett and Lemon, 1996]. Later, it was discovered that there are two kinds of synchrony among units – non-oscillatory and oscillatory [Baker et al., 1997]. Oscillatory synchrony is more closely related with LFPs and might be a plausible substrate for complex movements. In the motor cortex, the relationship between LFPs and unit activity has been observed to be quite complex. LFP oscillations are absent when actual movement occurs, while synchrony of units is still detected. Theoretical cortical models and experimental observations both suggest that LFPs are perhaps related more to general attention rather than related to a particular movement. LFP oscillations occur during the early phase of exploratory type of movements and are reduced when the monkey performed strongly trained movements [Murthy and Fetz, 1996]. LFP oscillations cease with trained tasks but do not occur with consistency during untrained movements, and most likely do not reflect neural discharge during such epochs [Donoghue et al., 1998]. Spikes also coincide with the phase of the LFP in certain frequency ranges during specific behavioral epochs [Murthy and Fetz, 1996], [Fries et al., 2001], [Womelsdorf et al., 2006]. Re-

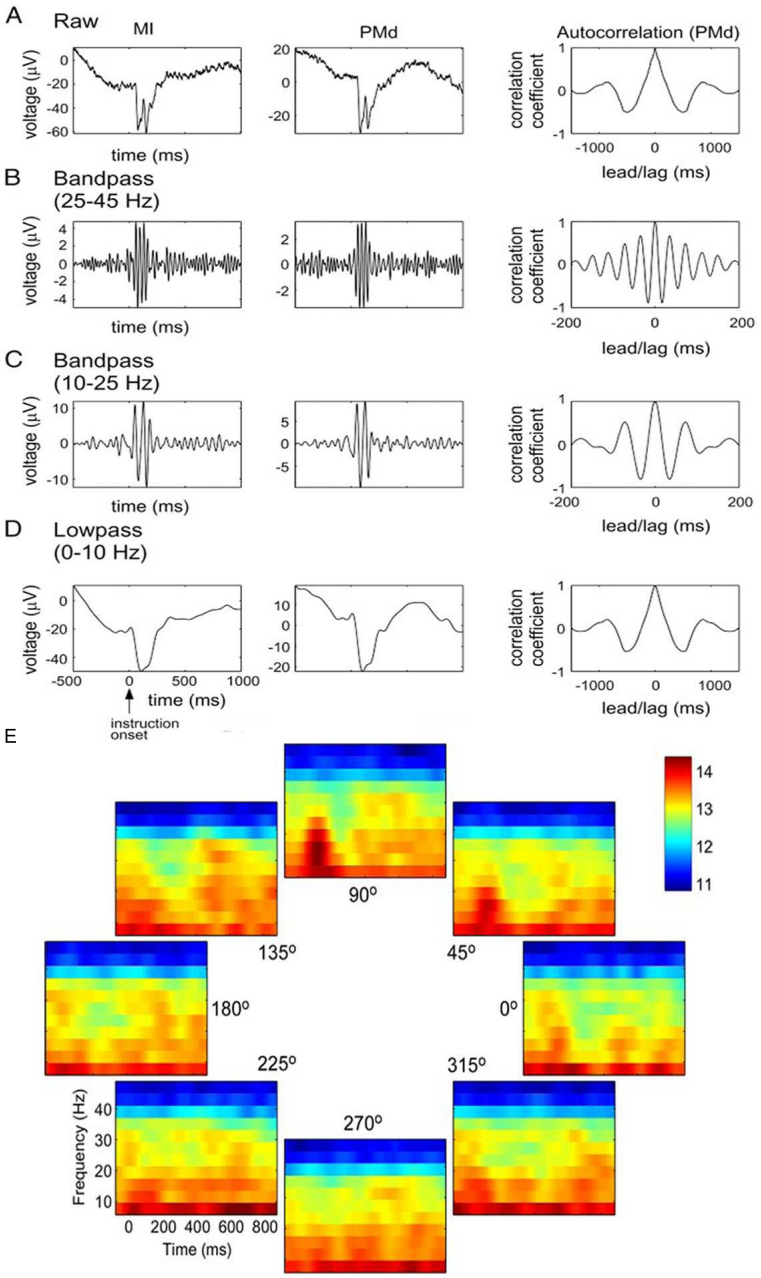

Figure 1.5: LFPs have a rich structure in both time and frequency domains and can be analyzed in time (as evoked potentials) with respect to behavioral events, or in different frequency bands. These approaches can also be combined as shown above. Figure shows LFPs recorded from two electrodes implanted in the primary motor cortex (MI) and dorsal premotor (PMd) cortex averaged over all trials, and the autocorrelation function of the evoked potential from the PMd A) raw LFPs. B) LFPs based on the average of band-pass (25-45Hz) filtered signals revealing phase-locked fast oscillations. C) LFPs based on the average of band-pass (10-25Hz) filtered signals revealing phase-locked intermediate (beta) oscillations. D) LFPs based on the average of lowpass (<10Hz) filtered signals revealing phase-locked slow fluctuations. E) Time-frequency spectrograms of the raw LFP from one channel recorded in PMd over all 8 target directions. For each target direction, single-trial spectrograms were computed and then averaged. There differences in the LFP activity with respect to movements to the different directions can be used for decoding. Activity in the 10-25Hz band was greatest in the 90° direction [O'Leary and Hatsopoulos, 2006].

cently, it was shown that correlations between spike and local field potential (LFP) activity between the PmD and the PRR are greater when monkeys are freely making choices compared to when they are following instructions [Pesaran et al., 2008]. Another study showed the LFP and unit activity recorded from the same electrode site showed different preferred directions [Asher et al., 2007].

All of the above studies suggest that LFPs and spike activity contain independent information about movement activity. Figure 1.5 shows the different temporal and spectral characteristics of LFPs recorded from the PmD and M1 for an 8-target task. Spectral changes have been observed in cortical surface potentials during motor movement [Miller et al., 2007]. As shown in Figure 1.6, different frequency ranges of local field potentials encoded different amounts of control signals and prediction power regarding the upcoming movement depended on the frequency band of choice [Rickert et al., 2005], [O'Leary and Hatsopoulos, 2006]. Decoding performance was the best in the lowest-frequency band (5-10Hz) and highest frequency bands (80-100Hz).

1.6 Justification and overview of dissertation experiments

In the context of the above mentioned body of work, we discuss the motivation for the studies in this dissertation. In all three studies, we use a rat model to investigate the the effectiveness of unit activity and local field potentials using chronic multisite silicon microelectrodes in the motor cortex.

Probabilistic tracking and assessment of neurons across sessions

While units can be recorded for a long-time with improved electrodes [Suner et al., 2005],[Vetter et al., 2004], stability of individual units is a challenge due to cortical pulsations and instability at the electrode-tissue interface

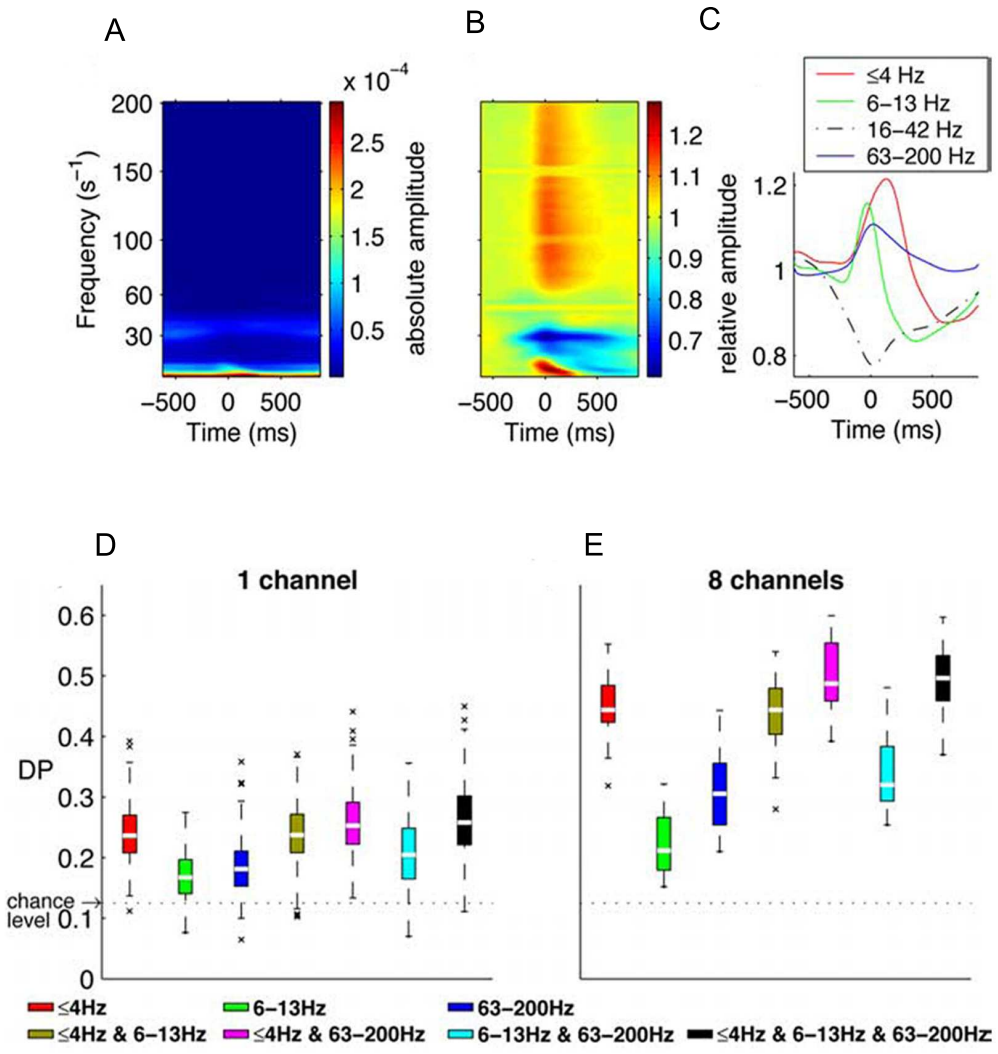

Figure 1.6: Time and frequency characteristics of LFPs in the $\pm 500ms$ period around movement onset (t=0): A) Trial-averaged LFP spectrogram showing the absolute amplitude; signal amplitudes are inversely proportional to the frequency. B) Time-resolved amplitude spectrum as in A), with each frequency bin normalized by its baseline amplitude shows peaks and valleys in different frequency bands around movement onset. C) Differences in evoked potentials in the different frequency bands (<4, 6-13, 16-42, and 63-200 Hz) during the task. D and E) Decoding power of different frequency bands for LFPs from a single channel and eight channels combined in the various frequency bands. The best decoding power was obtained by combining the lowest frequency bands(<4Hz and 6-13Hz) with the highest frequency band (63-200Hz) [Rickert et al., 2005].

[Patil and Turner, 2008]. Neuroprosthetic control algorithms require to be calibrated before every sessions. Being able to quantify neural stability from session to session will help in the design of neuroprosthetic algorithms and help in determining effective recalibration parameters [Andersen et al., 2004]. Techniques to reliably determine if the spike clusters across two different sessions represent the same neuron need further development. The goal while tracking neurons is to assess pairs of units from two different sessions and determine if they represent the same neuron based on the similarity of their neural signatures. There are two central assumptions: a) Neural similarity can be quantified by measuring some aspect of their neural signatures: shape, amplitude, negative-positive peaks, duration, or ISI, etc., (see Figure 1.7 B) and b) neural signatures recorded from neurons do not change abruptly, and on a short time-scale it is possible to track neurons from successive discontinuous sessions of recording by monitoring the changes in neural signature. The panel on the right in Figure 1.7 B shows the drift of cluster shape across sessions [Emondi et al., 2004].

The ideal approach to track single units would be to record from them continuously which allows monitoring of drifts and changes across time. Researchers have developed systems that allow such continuous recordings [Gilja et al., 2006], [Santhanam et al., 2007]; however, such recording systems may not be feasible for all experiments and instrumentation for such a system is far from trivial. Extracellular features of cells have been used to identify and classify types of cells [Barth et al., 2004],[Krimer et al., 2005]. There have been previous efforts to develop quantitative and automated techniques to track units from discontinuous recordings [Emondi et al., 2004],[Liu et al., 1999], [Liu et al., 2006], [Suner et al., 2005]. Figure 1.7 C demonstrates a recent approach to empirically determine the optimal separation by comparing the distribution of metrics comprised of different neurons with

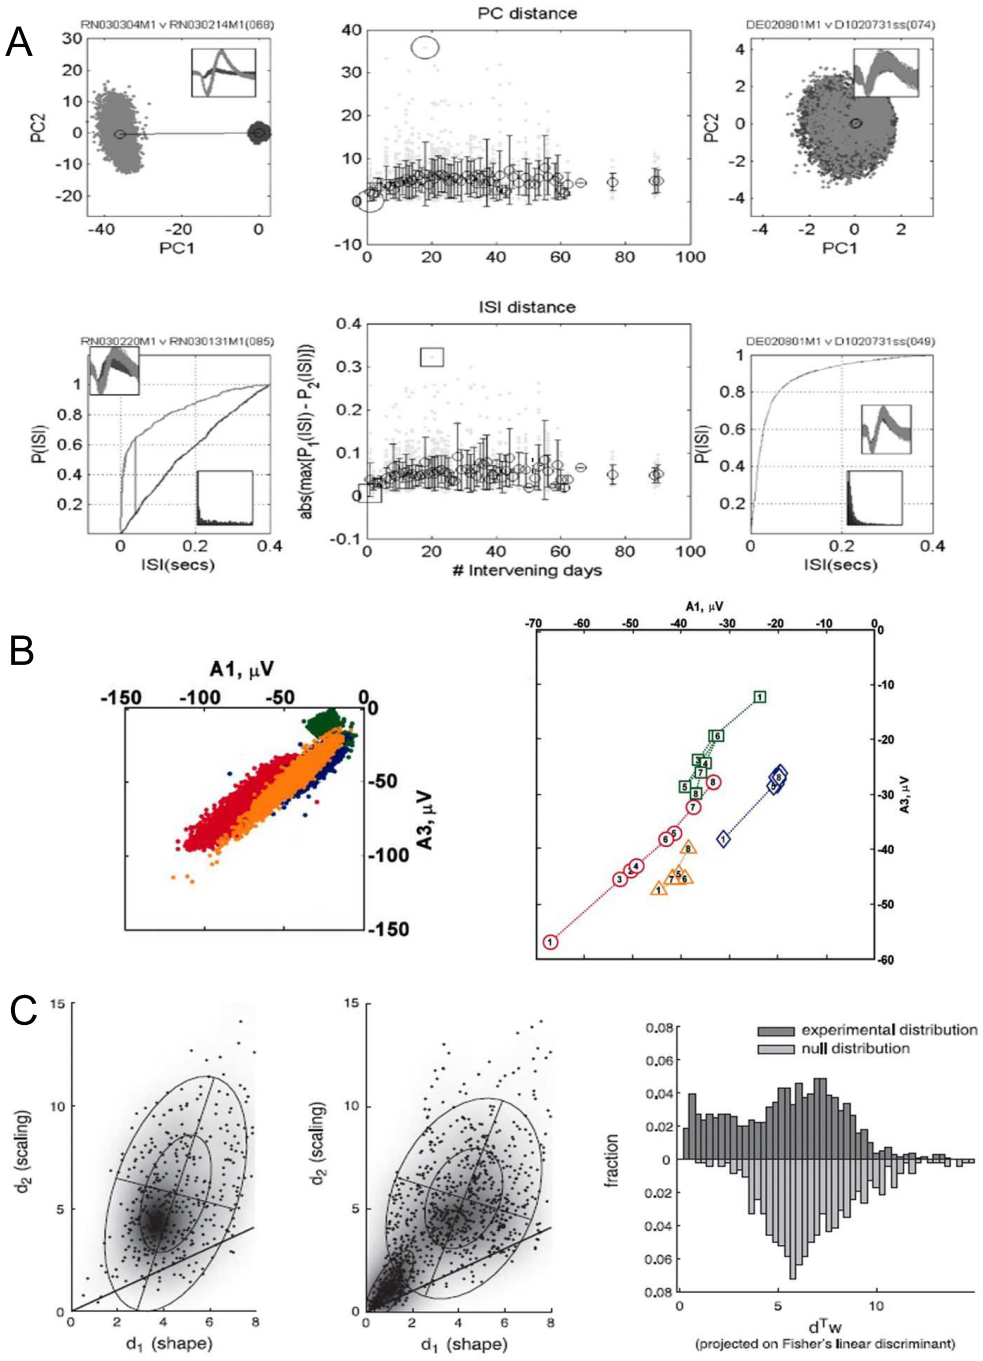

Figure 1.7: Different metrics used to track neural signatures and classify spike clusters. A) Middle panel shows metrics based on principal components (PCs) and inter-spike intervals (ISIs) across sessions of recording. Waveshapes, PC clusters, and the ISI distribution are shown for two sample sets of spike clusters that were classified as different (left panel) and the same (right panel) [Suner et al., 2005]. B) Left panel shows four sample spike clusters in 2-D amplitude feature space from one session of recording. Right panel shows the drift of the centroids of the four spike clusters in the feature space across sessions of recording [Emondi et al., 2004]. C) Another method to classify clusters is by comparing metric distributions of spike clusters from different neurons (null) and spike clusters under test (empirical). The distributions are unimodal (left panel) and bimodal (middle panel) respectively. Right panel shows the two distributions from the previous panels projected on a linear discriminant. The optimal separation boundary between the two modes to classify spike clusters as same or different in the empirical distribution is based on the null distribution [Tolias et al., 2007].

the distribution of metrics obtained from pairs under test [Tolias et al., 2007].

The goal of the first study is to present a novel technique that enables reliable tracking and assessment of neural signatures across different sessions for both single and multiple electrodes. We combined various previously published metrics under a logistic regression model that provides a probabilistic estimate of the 'similarity' between units obtained from two different sessions. This method has a number of applications including:

1. Neuroprosthetic algorithm testing by quantifying the stability of units recorded.
2. Validating the identity of a unit in operant conditioning while demonstrating effective neuroprosthetic control.
3. Analysis of cells in neural plasticity or learning experiments.
4. Quantifying changes in neural signatures across different sessions and assessing significance of particular metrics.

Units in the lower layers in the motor cortex are more effective targets for penetrating microelectrodes in cortical prostheses

A fast, reliable and fully-functional neuroprosthesis require development in better recording quality and higher stability from recording electrodes [Santhanam et al., 2006], [Patil and Turner, 2008]. While candidate neural sources and areas for control signals are being actively explored, the issue of how to best access these signals using implantable microelectrode arrays needs further study. Addition channels will surely improve performance, but there are constraints at the level of the neural interface resulting in cost-benefit tradeoffs. More number of electrode site counts add to the complexity of the packaging and multiplexing of electronics. Additional penetrating shanks may increase the incidence of deleterious reactive tissue response. One approach to tackle this issue is to consider interface designs that

are likely to yield the most salient control-related information with the fewest number of electrode sites and penetrating shanks. The goal is to maximize the information that can be extracted to control a neuroprosthetic device while minimizing tissue response and damage.

The neocortex exhibits a columnar structure consisting of six layers that is similar across the different areas [DeFelipe et al., 2002]. Can we exploit this structure for neuroprosthetic applications? Specifically, how does signal quality of recording in the different layers guide electrode design? How should electrode sites in the six-layered cortex be distributed to maximize salient control information? To the best of our knowledge there has been no systematic study to determine if the Layers 5,6 are indeed the best location for a neuroprosthetic device. In this study we investigate unit activity in the upper layers (2,3) versus the lower layers (5,6) in the motor cortex with respect to movement and direction. Locations of the electrode sites in the different layers are determined by using a combination of selective microlesioning and Nissl-staining of cortical slices for *post mortem* reconstruction of the electrode track. We also investigate ipsilateral or contralateral direction preference.

Laminar analysis of local field potentials (LFPs) in the motor cortex

From a neuroprosthetic standpoint, LFPs are suitable alternative control signals as they encode movement, direction, preparation and attentional features and are more stable than unit activity. Also, they can be simultaneously recorded from the same electrodes used to record unit activity. It has also been hypothesized that LFPs lack the spatial and temporal resolution of unit recordings [Donoghue et al., 1998], since LFPs are complex resultants of underlying synaptic activity and may not reflect the transient and specific synchronization of smaller groups of neurons. This implies that LFPs may reflect a global process active in conjunction with motor planning

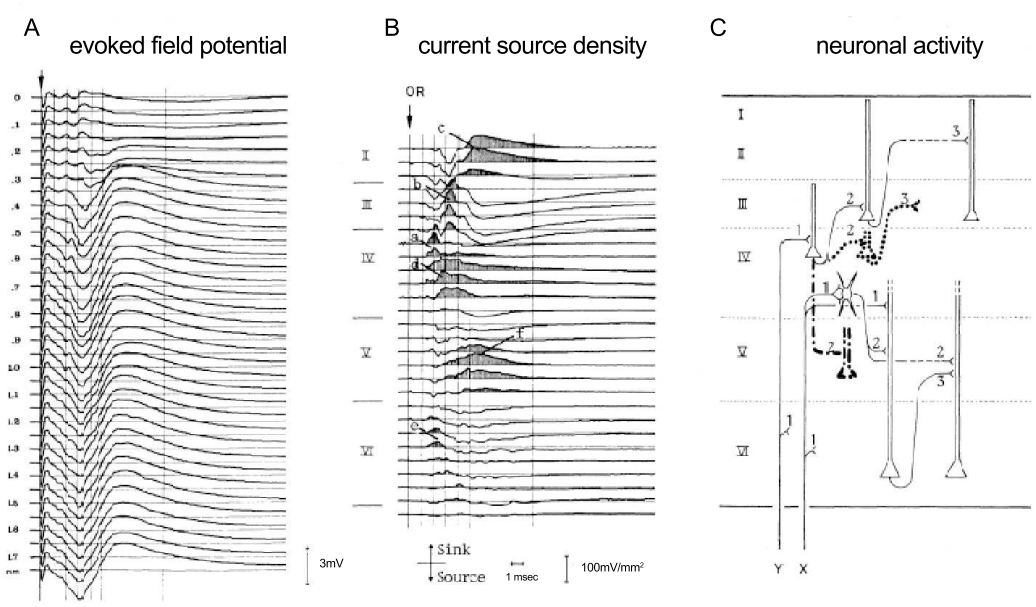

Figure 1.8: LFP activity varies across the different layers of the cortex: A) Evoked field potentials gradually change as a function of depth. B) Current-source density analysis is performed by taking a second-order spatial derivative of evoked potentials in A) to analyze current sources and sinks that give rise to field potentials; sinks and sources seen at different depths. C) Model of connections between cells in the different layers and neuronal activity to explain the sinks and sources that give rise to LFPs in the different layers of the cortex. [Mitzdorf, 1985]

or preparatory function, while the details of motor action are encoded in the unit activity.

Thus, it appears from previous studies that unit activity and LFPs encode different kinds of information about movement. Indeed, it has been shown that by combining the activity of LFPs and single units(SUAs) or LFPs and multi-units(MUAs), a superior prediction accuracy was obtained as compared to that provided by just using these signals alone [Mehring et al., 2003]. Signals conveying different cognitive functions are useful for creating multiple channels for communication [Andersen et al., 2004].

Apart from current-source density analysis, to our best knowledge, LFP activity has not been studied with respect to the lamina in the motor cortex. Given previous reports of the complex relationship of LFPs with spiking activity, we investigate if there are significant interactions in time and/or frequency across the different layers that can be exploited for neuroprosthetic use. The location of the electrodes in the different layers were determined by using a combination of selective microlesioning and Nissl-staining of cortical slices.

1.7 Dissertation Organization

The rest of the dissertation is organized as follows:

In Chapter II, we discuss a method to track and assess neurons across multiple sessions from chronic multisite electrodes. This manuscript is being prepared for submission to the *Journal of Neurophysiology: Innovative Methods*.

In Chapter III, we compare unit activity in the upper layers (2,3) and lower layers (5,6) with respect to movement and direction, and discuss the effectiveness of unit activity in the lower layers of motor cortex for neuroprosthetic applications. This

study has been accepted for publication in the *Journal of Neural Engineering*.

In Chapter IV, we investigate temporal and spectral aspects of LFPs across the layers of the cortex. This manuscript is being prepared for submission to the *Journal of Neural Engineering*.

In Chapter V, we conclude this dissertation by discussing the findings from the three studies in the context of the general field of neuroprosthetics, and suggest future work.

CHAPTER II

Tracking of Neurons

Abstract

We present a probabilistic framework to track and assess neural signatures across different sessions from both single and multiple electrodes. Being able to accurately identify and label spike clusters from one session to the next as the same neuron has important applications in modeling learning, studying neural plasticity, demonstrating effective neuroprosthetic control and designing algorithms. In other cases, a reliable method is needed to assess independence of spike clusters across different sessions to obtain an accurate sample size(n) for analysis. The common adopted approach, to track neurons across sessions, is to assess and quantify some neural signature, and then to determine a threshold, or a decision rule, to classify the spike clusters as the same neuron or different. What is absent from previous approaches is a quantitative assessment of the similarity between the spike clusters under consideration. Here, we demonstrate an automated method that incorporates different shape, amplitude, and timing metrics in one framework and provides a probabilistic measure of similarity. It can also be used as an additional method to quantify neural activity changes across sessions. Our analysis suggests that only a subset of metrics are required for a performance comparable to the full model. The timing metric (M_5)

that quantified inter-spike interval(ISI) differences was found to be a significant predictor of similarity. The method presented in this study is both flexible and scalable for different metrics and can be applied to both single electrodes and tetrodes.

2.1 Introduction

We introduce a framework using logistic regression that enables probabilistic assessment of similarity between spike clusters recorded across different sessions. In addition to classification, it also important to obtain a level of confidence in the classification. Studies modeling learning, understanding neural plasticity, in showing how a neuron alters its neural signatures under some behavioral or training paradigm [Marzullo et al., 2006], in assessing neural stability for neuroprosthetic control algorithms [Andersen et al., 2004]. We combined amplitude, shape and timing metrics that measure neural signatures in a single framework. This method also has applications in situations where data have been collected across multiple sessions, and we require a reliable method to determine accurate sample size(n) by excluding spike clusters that are the same from the different sessions.

With chronic multi-site electrodes it is now possible to record from an ensemble of neurons across multiple sessions for months, even up to a year [Suner et al., 2005]. Researchers have developed systems that allow continuous recording of unit activity [Gilja et al., 2006], [Santhanam et al., 2007]. Using such a recording system one can monitor small drifts and changes across time; however, such recording systems may not feasible as complexity of hardware is far from trivial and implementation of such a system may not be possible for many experiments. The attempt while tracking neurons is to assess pairs of spike clusters from two different discontinuous sessions and determine if they represent the same neuron based on the similarity of neural

signatures. Neural similarity can be quantified by measuring some aspect of their neural signatures: shape, amplitude, negative-positive peaks, duration, or inter-spike interval (ISI).

Visually labeling units using wave-shapes, ISIs, and auto-correlograms, etc., is an unreliable and time-consuming process. There have been previous efforts to develop quantitative and automated techniques to track units from discontinuous recordings [Emondi et al., 2004],[Liu et al., 1999], [Liu et al., 2006], [Suner et al., 2005]. All of the above approaches involve obtaining a metric that quantifies a particular neural signature: shape, amplitude, negative-positive peaks, duration, ISI, etc. The neural signatures, or corresponding metrics, are compared and then some threshold criterion applied to determine if the spike clusters are from the the same neuron. At the same time, similarity of neural signatures does not prove that they are from the same neuron. Hence, a recent approach was to compare the empirically obtained distributions of metrics from spike clusters that were known to be from different neurons (null distribution) with the distribution of metrics obtained from unit pairs that were being tracked (test distribution). The optimal separation was then determined by comparing the null distribution with the test distribution [Tolias et al., 2007]. In addition to classification, a further improvement would be to have a probabilistic measure that assigns a level of confidence in the similarity[Emondi et al., 2004]. This would help assess the closeness of spike clusters in metric space and specify a tolerance for separation.

Our attempts to directly adapt the method for tracking neurons from tetrode data as proposed by [Tolias et al., 2007] to our single electrode data were not successful (our unpublished observations). We concluded that, at least for our single electrode dataset, using just the two metrics proposed was not sufficient to classify

and track units. In comparison to single electrodes, tetrode recordings contain more information about the neuron as they record the neural waveform from four spatial positions. Thus, for single electrode recordings, using additional metrics may help in better classification. None of the previous methods have analyzed the different neural signatures in one framework and determined their effectiveness in tracking.

In this study, we evaluate a number of different metrics in a single framework and show an automated method to track and assess neural signatures obtained from chronic multisite recordings across sessions from both single electrodes and tetrodes. The novelty and utility of using logistic regression is that it provides a probabilistic estimate of the similarity or 'closeness' between the two spike clusters under consideration. A probabilistic method will enable additional quantification of electrode drift and changes due to neural plasticity. Also, depending on the application or hypothesis under test, the decision criterion for classification can be adjusted to be more or less conservative. This method can be used for data recorded from both single and multiple electrodes. Depending on the area of the cortex, type of cells recorded, or nature of the experiment, the method can be tuned to emphasize particular metrics.

2.2 Methods

Tetrode and single electrode data was used for training, testing and verification of the technique. The tetrode data was collected from four rats and was sorted manually. Six tetrode channels of data were recorded from the medioprefrontal cortex (infralimbic and prelimbic), and twelve additional tetrode channels were recorded from the amygdala (lateral, basolateral, and central). Surgical methods and data collection procedures were similar to those described in [Berke et al., 2004]. Single electrode data was collected from multi-site recordings, each from a single sixteen-

site electrode array (NeuroNexus Technologies, Ann Arbor, MI) implanted in the motor cortex of three Long-Evans rats (D2-D4). Details of the surgical procedure and location of the probes have been described in [Parikh et al., 2009].

2.2.1 Description of metrics used

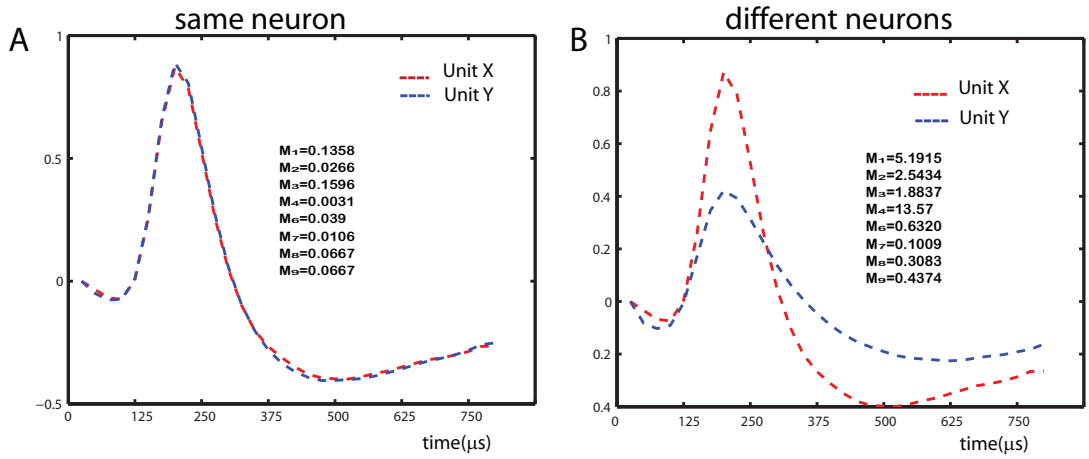

Figure 2.1: Mean waveforms, amplitude and shape metrics: Each panel shows the mean waveforms and measured metrics for typical spike clusters — Unit X and Unit Y — obtained from the (A) same neuron and (B) from two different neurons. All metrics M_1-M_9 (except M_5) quantify some aspect of amplitude and/or waveshape differences between the spike clusters under test. As can be expected, the metrics for the same neuron are much smaller than metrics obtained by comparing two different neurons.

We used, or adapted, previously published metrics to compare neural signatures between the pairs of units under test as described below. It should be noted that these metrics are not necessarily independent measures of neural signatures. Most metrics quantify some aspect of waveform or amplitude or both. One objective of this study was to determine the most effective metrics. Although, we only considered a few metrics, additional metrics can be easily incorporated in the subsequent regression model. All the metrics presented here are pairwise metrics, and give the relative distance between the two spike clusters consideration: Unit 'X' and Unit 'Y' Thus, for spike clusters from the same neuron we expect the resulting pairwise metrics to be ideally at, or very close to, the origin in multi-dimensional metric space.

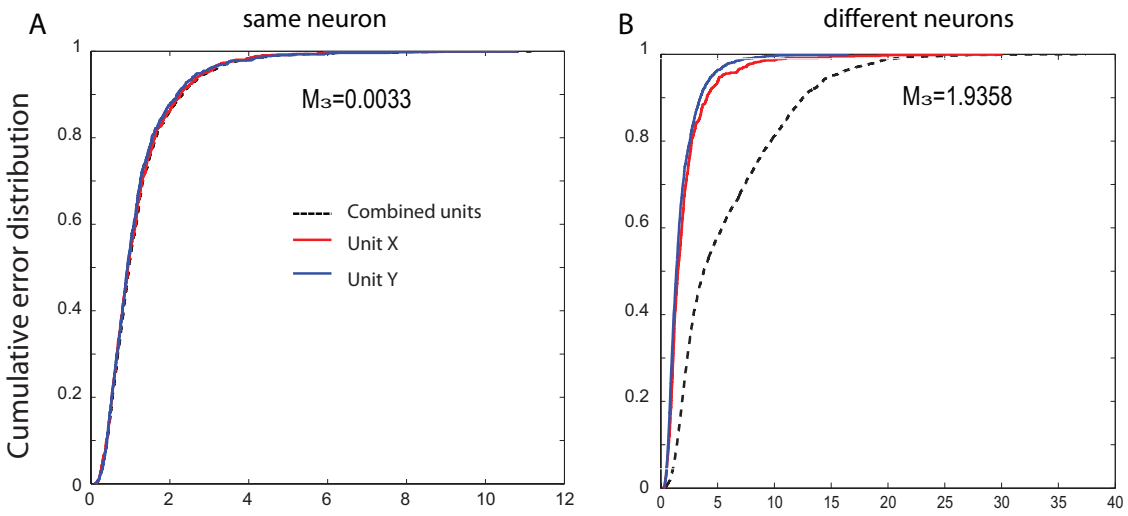

Figure 2.2: Cumulative error metric (M_3): Squared difference, or error, between each individual waveform and the mean waveform was calculated. Each panel shows the cumulative error distribution for errors obtained after combining the two spike clusters (shown in black), and the cumulative error distribution for Unit X (red) and Unit Y (blue) considered individually, from the (A) same neuron and (B) from a different neuron. As expected, the differences in the cumulative error distributions for the individual units and the combined units from the same neuron are smaller than those for units that represent two different neurons.

M_1 and M_2: These two metrics measure the difference in amplitude and shape of the mean waveforms of Unit X and Y as originally described in [Tolias et al., 2007]. For a tetrode there are n=4 channels or contacts, and for a single electrode, n=1. For each of the 'n' electrode channels, the average waveform of Unit X is scaled to minimize the sum of squared differences between the average waveforms of Unit X and Y; the scaling factor of X and Y is denoted as $\alpha(X_i, Y_i)$, where X_i and Y_i are vectors of length equal to the digitized samples of the thresholded action potential. This solely captures the difference in shape because the Unit X has been scaled to match Unit Y, and both have further been scaled by Unit Y. d_1 is a normalized Euclidean distance between the scaled waveforms and is computed as,

$$d_1(X,Y) = \sum_{i=1}^{n} \frac{\|\alpha(X_i, Y_i)X_i - Y_i\|}{\|Y_i\|} \tag{2.1}$$

where, n=no. of channels, e.g., 4 for a tetrode. The difference in amplitude was

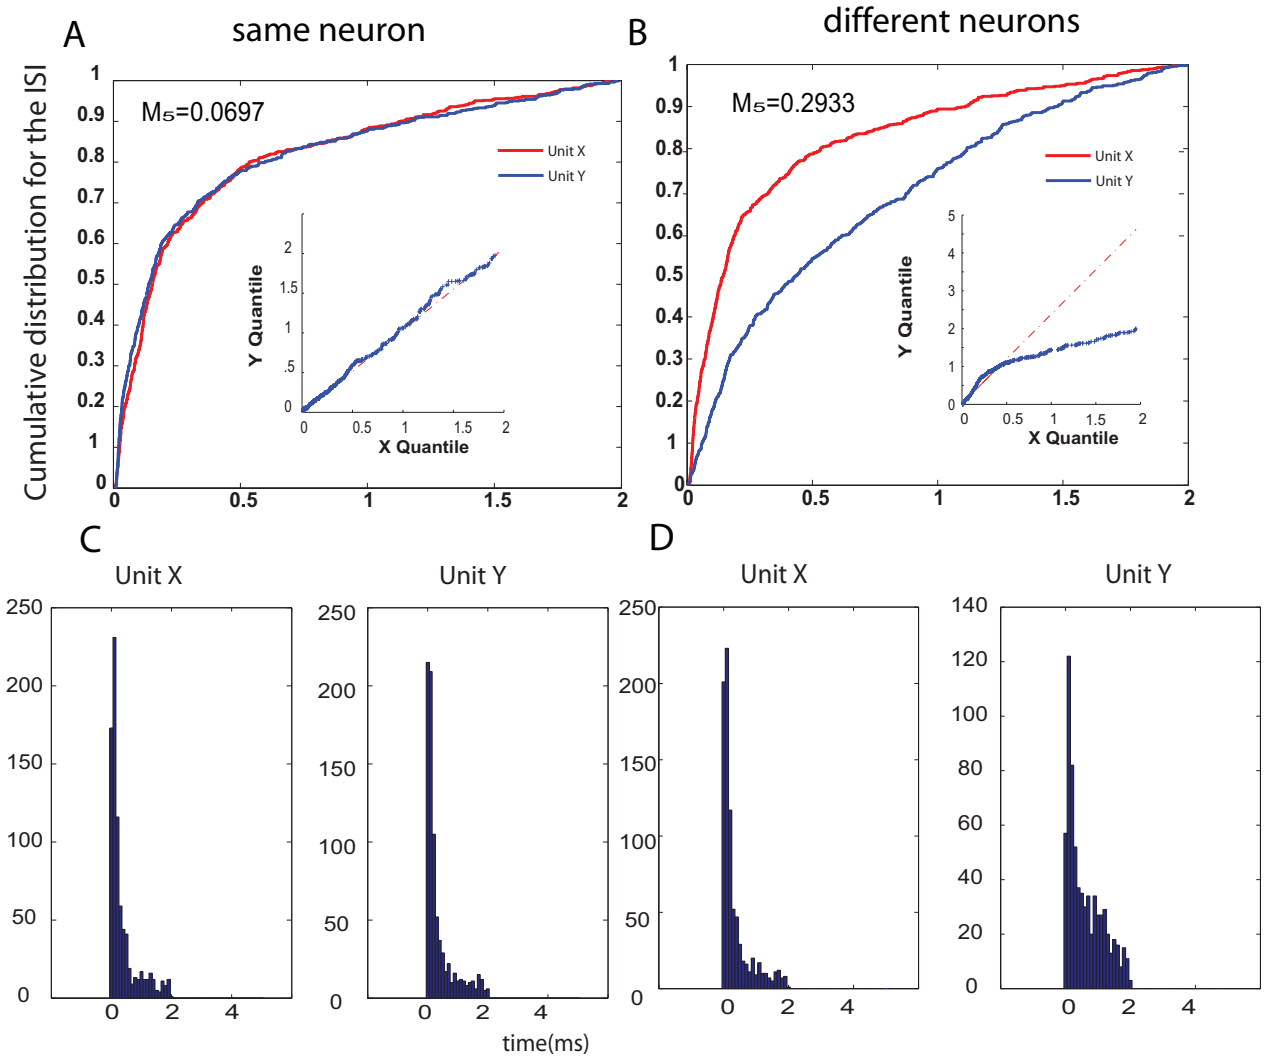

Figure 2.3: Timing metric (M_5): Top panel shows the cumulative distribution of the interspike intervals (ISIs) for two sample units X and Y that are obtained from the (A) same neuron and from two (B) different neurons. Inset panel show quantile-quantile(q-q) plots for the ISI distributions. For similar distributions, the quantiles for both units should lie along the 45° line in the q-q plot. Bottom panel shows the histogram of the ISI distribution for both units in the two cases.

quantified by d_2 computed using the scaling factors calculated earlier for each of the 'n' channels denoted by 'i' and 'j',

$$d_2(X,Y) = \max_{i=1}^{n} |\ln \alpha(X_i, Y_i)| + \max_{i,j}^{n} |\ln \alpha(X_i, Y_i) - \ln \alpha(X_j, Y_j)| \tag{2.2}$$

The final metrics are given by: $M_1 = d_1(X,Y) + d_1(Y,X)$, and $M_2 = d_2(X,Y) + d_2(Y,X)$ for symmetry. Figure 2.1 show examples of two spike clusters X and Y that are from the same neuron and different neurons.

M_3: This metric measures the variance of the individual waveforms. First, we compute all the squared errors of the individual waveforms from the mean waveforms when the two units X and Y are considered separately; second, we compare the squared errors with the mean waveform obtained after combining the two units. The errors are calculated as,

$$\delta = \sum_{i=1}^{n} (X_i - \overline{X})^2 \tag{2.3}$$

Then we compute the Kolmogorov-Smirnov (KS) statistic for the cumulative error distributions for Unit X and Unit Y with the δ_{comb} distribution for the combined units, i.e. X+Y. The chief advantage of the KS statistic is that it makes no assumption on the distribution of the two cumulative distribution functions (cdfs) under consideration. If the two units are from different neurons, their mean waveforms are different and will result in a larger KS statistic with respect to the combined waveform. M_3 is sum of the two KS-statistics obtained after comparing distributions of X and Y with respect to the combined unit. Figure 2.2 shows that for units that are from the same neuron, the individual and combined error distributions are approximately the same, resulting in a smaller value of M_3 relative to the value obtained when two different neurons are compared.

M_4: This metric measures the Euclidean distance between the mean waveforms of Unit X and Y which are 'N'-dimensional vectors, where N corresponds to the number of digitized samples of the waveform. In our case, Unit X and Unit Y were vectors of length N=32. M_4 was computed as,

$$M_4 = \sum_{j=1}^{n} \sqrt{\sum_{i=1}^{N} (Y_{j,i} - X_{j,i})^2}, \tag{2.4}$$

where 'n' corresponds to the number of contacts, for tetrodes: n=4. If the mean waveforms are the same, then the norm or the Euclidean distance between them will be closer to the origin as compared to that obtained from two different neurons. For example: Figure 2.1 shows the mean waveforms for the same unit which have M_4=0.0031, compared to the M_4=13.57 for mean waveforms from the two different units. In comparison to M_1 which only considers shape differences, M_4 considers both shape and amplitude differences, as difference in either aspect would results in a larger m_4 metric.

M_5: In addition to waveshape, the timing of the firing is an important statistic. The interspike interval (ISI) is often used to classify neurons. As shown in Figure 2.3, the ISIs are different for the same and different neurons. M_5 measures the two-sample Kolmogorov-Smirnov test statistic (KS Test) between the cumulative distribution of the ISIs of the two units under test as shown in Figure 2.3. The quantile-quantile (q-q) plot compared the distributions of the ISIs. For the same neuron, the estimated quantiles should be equal and should lie on the 45° line. The assumptions in using this metric are that there are no state changes (eg: awake vs. anaesthetized) and no abrupt changes in the ISI due to learning between sessions under comparison.

M_6-M_9: This set of metrics measures the amplitude and shape differences between the two units under test based on the amplitude of the peaks and duration of the

positive and negative phases of the mean waveform as described in [Liu et al., 2006]. We modified the metrics to enable relative comparison and for the metrics to be in the range $[0, \infty)$. M_6 compares the normalized relative ratio of the peak-to-peak amplitude(A_{pp}) of the mean waveforms X and Y given by the formula:

$$M_6 = \| \frac{A_X - A_Y}{A_X + A_Y} \| \tag{2.5}$$

M_7 compares the normalized relative ratios between the negative-to-positive (NP) peaks of the waveforms X, Y under test.

$$M_7 = \| \frac{NP_X - NP_Y}{NP_X + NP_Y} \| \tag{2.6}$$

Similarly, M_8 and M_9 compare the ratios of the negative phase and ratios of the negative to the positive phase duration respectively.

2.2.2 Modeling using Logistic Regression

Given the set of predictors in Section 2.2.1, the problem is to predict whether the spike clusters are from the same or different neurons. Here the response (denoted henceforth by S) is binary. S=1 indicates the same neuron, and S=0 indicates units are from different neurons. Binary logistic regression is a natural choice in this case as it gives the probability of belonging to one class or the other. Given the set of predictors $M = M_1, ..., M_p$, we denote, $P(S = 1|M) = p(x)$. Then the model is described by,

$$\log \frac{p(x)}{1 - p(x)} = \beta_0 + \beta_1 M_1 + \ldots + \beta_p M_p + \eta \tag{2.7}$$

where, $\eta \sim N(0, \sigma^2)$ is the normally distributed error term. Logistic regression provides an estimate of conditional probability of being in a particular class $P(S = 1|M)$,

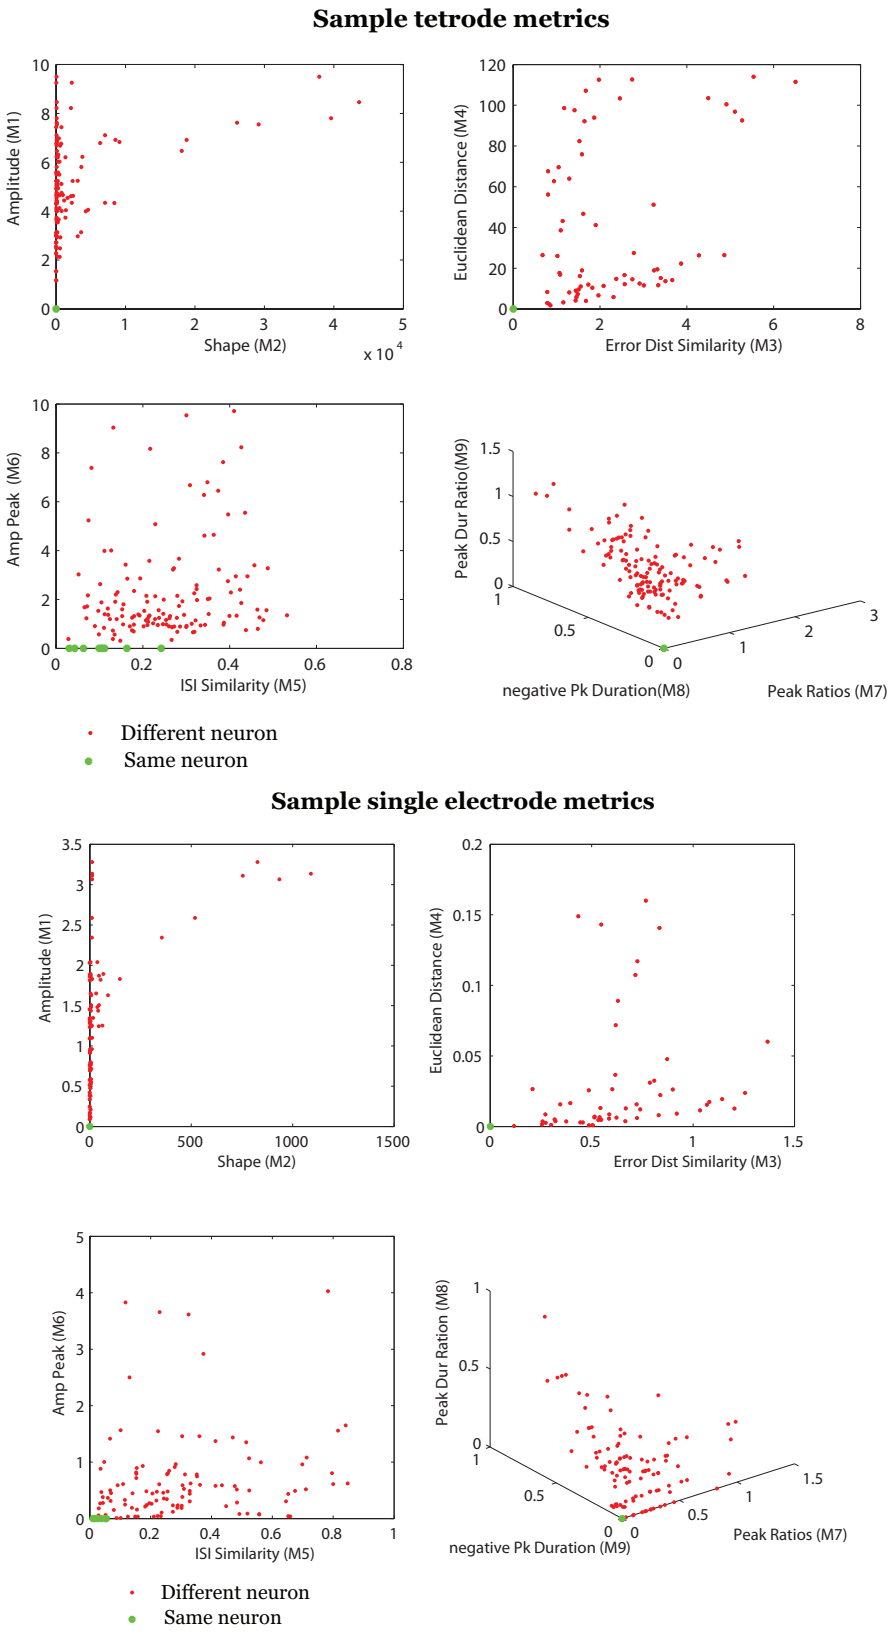

Figure 2.4: Sample metrics for tetrode and single site locations: Panels plot the metrics M_1-M_9, shown in two-dimensional and three-dimensional metric space for visualization purposes. Data show that pairwise-metrics from the same neuron are clustered at the origin (shown in green), and pairwise-metrics for different neurons (shown in red) are distributed away from the origin.

given the set of predictors. As an initial classification, the unit pair is considered to be the same if the probability is closer to 1.0 than 0. Hence, the default decision boundary is set at 0.5. More generally, this decision boundary can be moved from $p(x) = 0.5$ based on the desired false positive (classifying spike clusters as the same, when they are not) or false negative rate (classifying a spike clusters as the different, when they are from the same neuron). For example: if we can be more conservative in our false positive rate, as move the decision boundary close to 1. Two separate decision boundaries can be used to be more conservative with regards to false positives and false negatives at different levels of confidence. If a computed probability falls between the two decision boundaries, then we classify the spike clusters as 'uncertain'.

We implemented the logistic regression using the `mnrfit` routine in MATLAB (Mathworks, Natick, MA), and further verified with the `glm` routine using the R statistical package (http://www.r-project.org/). All computed pairwise metrics (M_1-M_9) for the spike clusters being compared were normalized by subtracting the mean and dividing by the standard deviation before logistic regression was carried out. In many cases, the residual deviance of the fit is extremely small indicating a very good fit, but the predictors are not significant due to high standard errors. This often arises in datasets where the groups are linearly separable - resulting in a perfect fit, but unstable parameter estimates for the same reason [Faraway, 2006]. The instability arises from the fact that multiple parameters can achieve a similar separation due to the highly separable data. In particular, the maximum likelihood estimates of the model parameters are biased away from 0. In such a case, we need to use bias reduction method due to [Firth, 1993]. This enables us to remove the $O(n^{-1})$ term from the maximum likelihood estimate and gives a model fit that is much more

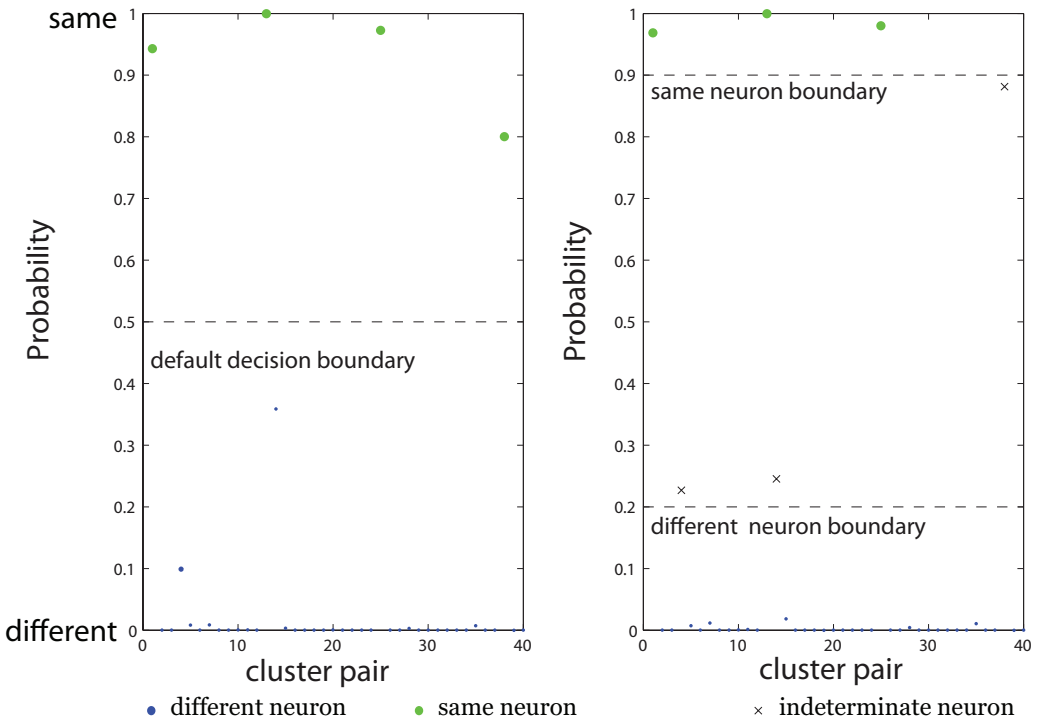

Figure 2.5: Sample bounds: The decision bounds can be changed depending on level of confidence required in the classification. In the first panel on left, neurons are classified as the same (shown in green), or classified as different units (shown in blue) if their predicted probability was greater or lower than 0.5 respectively. In the second panel, the bounds are made much stricter: the upper bound is at 0.9 and the lower bound is at 0.2. Spike cluster pairs lying in between these regions are marked as indeterminate.

stable for prediction purposes. The bias reduction method was implemented using the `brglm` library in R.

2.2.3 Obtaining model parameters and predicting responses

First, we need to train the model and obtain the regression parameters with known responses (S), before we can predict similar units across two sessions. For one of the session being compared, we created two pseudo-sessions: 'A' and 'B'. To avoid a training bias, the pseudo-sessions were created by shuffling the order of the individual samples in the original session and then assigning them A or B. We compute metrics between spike clusters in the two pseudo-sessions and set the response to S=1 or 0, depending on whether it was the same unit or not. Once the model parameters are obtained, we compare spike clusters in the two different sessions and compute the probability of similarity between the spike clusters using the regression coefficients obtained. The default decision threshold was set at 0.5. This decision threshold can be shifted, or two decision thresholds used depending on the application and level of confidence required in the assessment of similarity between two clusters.

2.2.4 Cross Validation and testing of the method

The prediction performance was tested by cross validation using a Monte Carlo method. Since parameter estimates are dependent on the training data, we wished to determine the stability and robustness when the training data was varied. Thus, in addition to the pseudo-sessions 'A' and 'B', we created an additional pseudo-session for validation called 'Test' using the same shuffling procedure as described earlier. The model was fitted as described and the prediction validated using the 'Test' data. Bootstrap samples were drawn repeatedly from the original session to create the

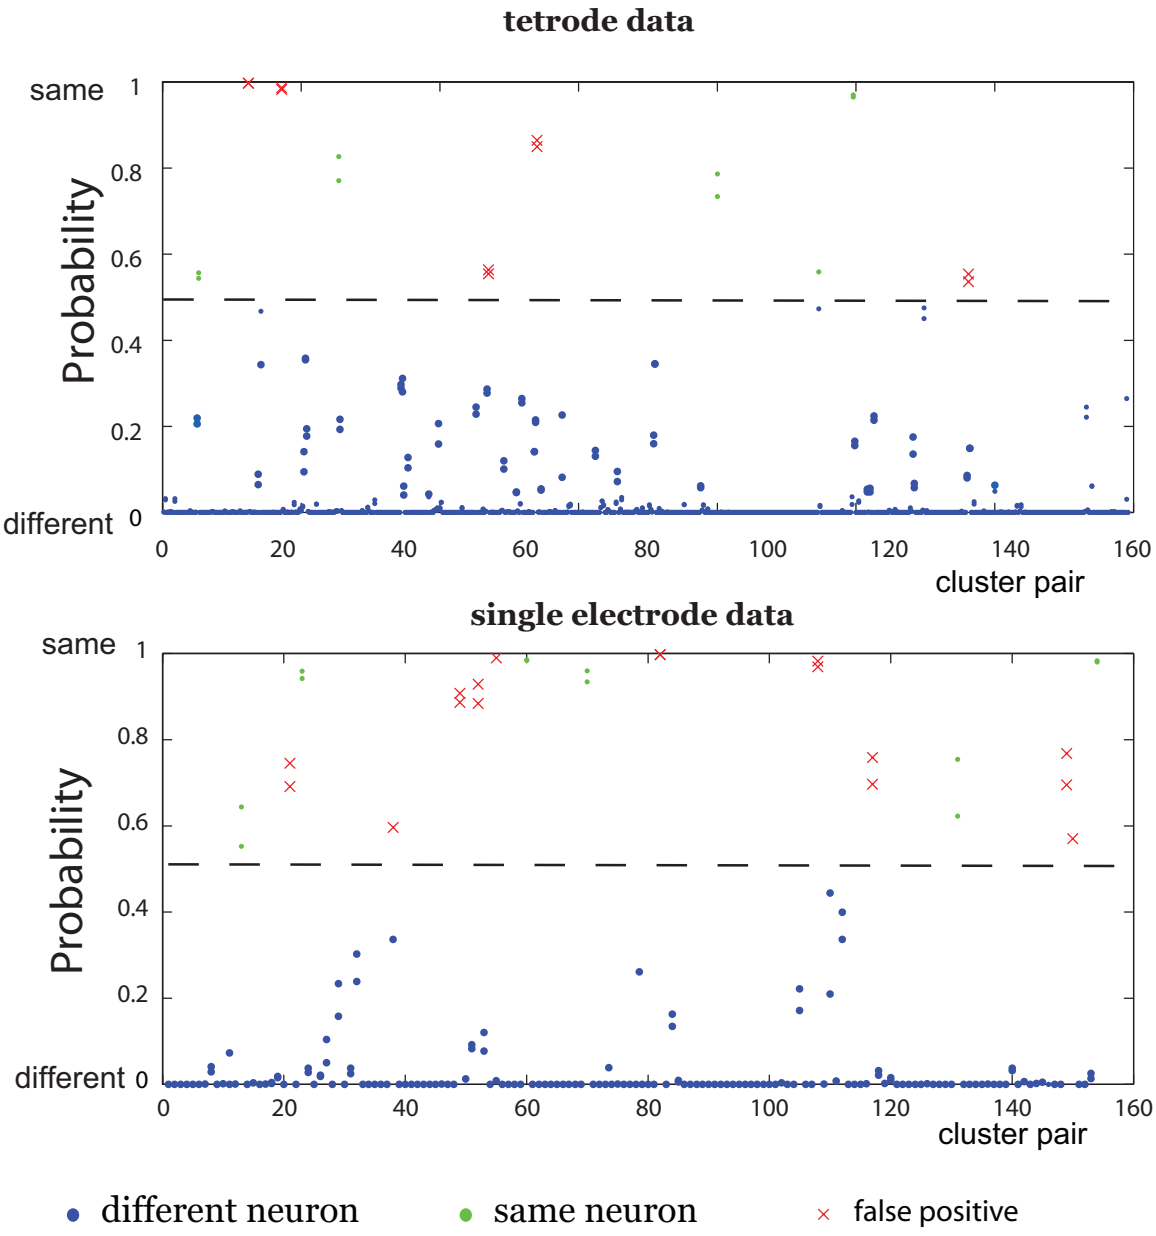

Figure 2.6: Predictions for tetrode and single site locations: Panel shows performance for two sample datasets for tetrode and single electrode data. The decision boundary was set at a default of 0.5. Spike clusters with a probability greater than 0.5 were classified as the same (shown in green) or different (in blue). A few unit pairs (7-9%) from completely different channels were predicted to have a probability >0.5, these false positives (red crosses). Performance for single and tetrode data was comparable.

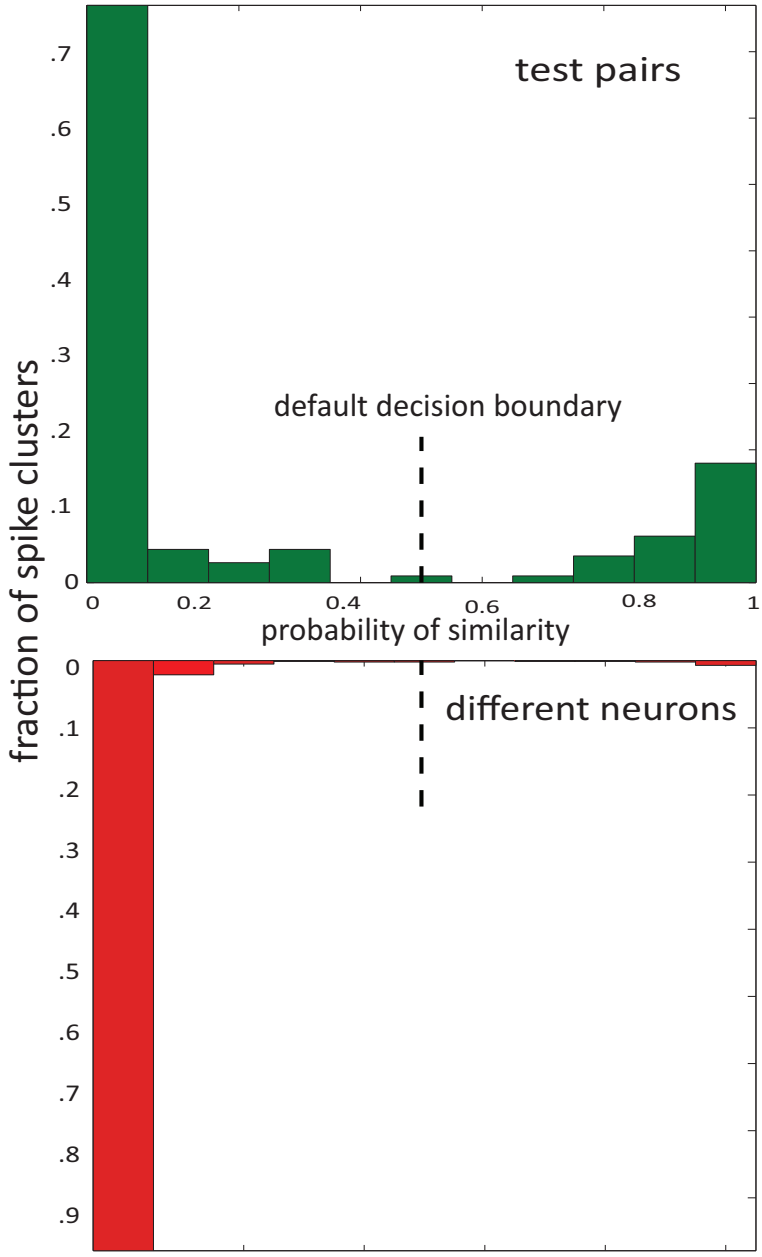

Figure 2.7: Distribution of computed probabilities for pairs from the same channel (test pairs) and from pairs of neurons from different channels. Data show probability of similarity distributions for spike cluster pairs obtained from 9 successive sessions. The test pair distribution is bimodal and distribution for different neurons is unimodal. The overlap between the two distributions is minimal. The decision boundary for classification can be adjusted to obtain a better level of confidence depending on the application.

three pseudo-sessions and performance statistics obtained for 1,000 such iterations of the above process.

2.2.5 Comparison with behavioral correlates

It is expected, that spike clusters that are identified as the same across two sessions should encode the same behavioral parameters, such as encoding the same place field, movement to a particular direction, orientation tuning function, etc. Similarly, spike clusters identified as different should on average have more dissimilar behavioral correlates. As an additional method to verify our results we compared behavioral correlates for spike clusters identified as the same across two different sessions. This single electrode data was collected from animals that were trained on a two-direction movement discrimination task as described in the subsequent ChapterIII. We analyzed 14 pairs of sessions from six implants and compared firing rate responses in analysis epochs as shown in Figure 3.1. We compared three behavioral correlates, a) for movement encoding, a unit either showed a significant relative increase, decrease, or no significant change when the firing rate in the 500ms post-movement epoch was compared to the 500ms pre-movement epoch, b) and c) In two different analysis windows we compared direction encoding by comparing firing rate responses in rightward and leftward trials. A unit encoded directional information if it either showed a significant differential response towards leftward or rightward movements.

The behavioral correlate coefficient ($R_{x,y}$) between all unit pairs (x,y) was computed as,

$$R_{x,y} = \sum_{i=1}^{N} \frac{m_i}{N} \tag{2.8}$$

where, $m_i = 1$ denotes a behavioral match for the i^{th} behavioral correlate, and correspondingly, $m_i = 0$ for a non-match; for our single electrode data, we compared

three behavioral correlates hence, N=3 .

Thus, $R_{x,y}$ is '1' when all correlates perfectly match, and '0' when there are no behavioral matches. Since the units being compared are located in the same region of the brain, it is expected that spike clusters that are detected as different may still have the same behavioral correlates, or these matches occur due to chance. However, we expect that the distribution for spike clusters that were detected the same (x=y) should be different from the distribution for spike clusters that were detected to be different (x $\neq$ y).

2.3 Results

2.3.1 Logistic regression analysis

Figure 2.4 shows each of the different metrics (M_1 to M_9) with the same and different units from the training datasets plotted in 'green' and 'blue' respectively. The parameters obtained from the pseudo-sessions were then used to predict the probabilities of the spike clusters from the test dataset. Figure 2.5 shows typical results obtained from applying our method to the pseudo-sessions. Neurons classified as the same are shown in green, and different neurons are shown in blue. There were a few misclassifications: some pairs from different channels were predicted to have a probability below the decision threshold by the model. In practice, these false positives can be discarded. Figure 2.6 shows a sample prediction for two actual datasets. As expected, for real datasets, the errors are slightly higher than those predicted from the pseudo-session Monte-Carlo analysis, due to temporal drifts in the metrics. Figure 2.7 shows the distribution of the probabilities for different neurons and spike clusters from the same electrode site. The distribution of the different neurons is unimodal, and the bimodal for the test pairs indicating that neurons on the same electrode channel may be either different or the same.

Metric	M_1	M_2	M_3	M_4	M_5	M_6	M_7	M_8	M_9
M_1	1.00	0.50	0.47	0.36	0.22	0.51	**0.79**	0.12	0.25
M_2		1.00	**0.74**	**0.88**	0.28	**0.99**	0.30	0.02	0.14
M_3			1.00	0.59	0.12	**0.77**	0.31	0.13	0.17
M_4				1.00	0.16	**0.87**	0.08	-0.02	0.06
M_5					1.00	0.26	0.26	0.02	0.10
M_6						1.00	0.30	0.02	0.14
M_7							1.00	0.09	0.17
M_8								1.00	**0.69**
M_9									1.00

Metric	M_1	M_2	M_3	M_4	M_5	M_6	M_7	M_8	M_9	remarks
Full Model	0.4900	0.7878	0.4347	0.6202	0.0319 *	0.3172	0.8608	0.5148	0.7101	AIC: 173.22
Reduced Model	0.3098	0.0149 *	–	–	0.0062 **	–	–	–	–	AIC: 163.65

Table 2.1: Metric evaluation: The first table shows a sample metric cross-correlation plot. Metrics are not independent; all metrics may not be required while fitting the logistic regression model. Correlations above 0.6 have been shown in bold. M_3 and M_6 were are often found to be correlated to the other metrics. The timing metric M_5 showed lower correlations to the other metrics. The second table shows the significance of the different predictors under a full model and a reduced model for a sample dataset and corresponding Akaike Information Criterion (AIC) values. Prediction performance was not comprised with a reduced model.

In many of our datasets, ordinary logistic regression did not converge, as the two groups were linearly separable. Due to the high separability, we can fit the data perfectly, resulting in extremely low values of residual deviance. We can obtain an excellent model fit, but none of the predictors are significant individually and the standard error of the estimates were also extremely high. For such cases, as recommended [Firth, 1993], we used the bias reduction method in order to get stable parameter estimates. We also analyzed the significance and effect of different metrics on the predictive performance. We used the AIC to obtain the most parsimonious model.

2.3.2 Cross-correlation and significance analysis of the metrics

All metrics that we tested are not necessarily independent measures of neural signatures. Most metrics quantified some aspect of waveform, amplitude, or both. Table 2.1 shows sample correlation values between the metrics. Generally, across the different datasets, we found that M_1 (amplitude) is strongly correlated with M_6 or M_7. This result was expected as these three metrics measure amplitude character-

Upper Bound	Total pairs tested	correct classifications	false positives	Indeterminate
.51	2000	1798	2	0
.6	2000	1796	1	22
.7	2000	1789	0	28
.8	2000	1788	0	23
.9	2000	1747	0	66

Table 2.2: Monte Carlo evaluation of performance for a sample dataset pair. Statistics for 1,000 repetitions under different upper bound conditions. False positive rates fall as the bounds are made tighter. Correct classifications decrease as a function of a tighter upper bound, as more predictions fall in the indeterminate range.

istics. We also determined that M_3 (cumulative errors) was often correlated to the other metrics, and could be discarded in the regression. When the significance of predictors was taken in consideration, M_8 and M_9 were rarely found to be significant and were often correlated to each other. Hence M_3, M_8 and M_9 were often dropped to obtain a reduced model with a lower Akaike Information Criterion (AIC). M_5, the timing, or ISI distribution metric, showed low correlations with the other metrics, and was found to be significant predictor in the logistic regression in many of the session analyzed. This suggests that incorporation of timing information in addition to static characteristics, such as waveshape and amplitude, can help in better classification and tracking. Prediction performance remained virtually unchanged with a reduced model.

2.3.3 Effect of changing the decision boundary on performance

To determine the effect of changing the decision boundary, Monte-Carlo cross-validation was performed. Table 2.2 shows typical results for repeated tests for one session pair. The lower bound was kept fixed at 0.3, and we systematically varied the upper bound. The table shows that as the upper bound was made tighter, i.e. closer to 1.0, the number of detections decreased. Correspondingly, the number of unit pairs falling in the indeterminate case increased and false positives also decreased. In a separate analysis, we analyzed the effect of how the choice of the lower bound

affects the number of units that are classified as different, including the ones from the same channel which were detected to be different. Using the pseudo-session data, we evaluated the performance of false negatives, i.e. units that were the same but were detected as different, which is a more critical test of performance. Using the Monte Carlo process, we found that the error rates were less than 1%.

False positive errors are less critical and they can be easily discarded since we know that in these cases the spikes clusters were recorded from electrode locations that were physically distant and hence cannot possibly be the same. There are a few explanations for obtaining false positives: 1) These could arise purely by chance since we are making so many comparisons, 2) Pairs of neurons may be distinct but could share the same characteristics, since we are recorded from the same region of the brain.

2.3.4 Behavioral Correlates

We compared the behavioral correlate coefficient ($R_{x,y}$) for clusters identified as the same (i.e. x=y) with the behavioral correlates ($R_{x,y}$) with all other units (where, x $\neq$ y). Figure 2.8 shows the fraction of the behavioral correlates from three animals. The Wilcoxon ranksum test for the data showed that the two distributions obtained are different (p=1.05×10^{-8},z= 5.72, ranksum: 2.824×10^4, $n_1 = 56$, $n_2 = 661$). Some of overlap between the two distributions can be expected by chance. Overlap between the two distributions can also be expected for the following reasons: a) Since neurons are recorded from the same location in the cortex, different neurons are likely to have similar tuning functions. b) If the electrode moves slightly due to micromotion, the neural signature will change to be detected as different, but since it is the same neuron it will have the same behavioral correlates [Tolias et al., 2007]. An explanation for observing the same unit with different behavioral correlation

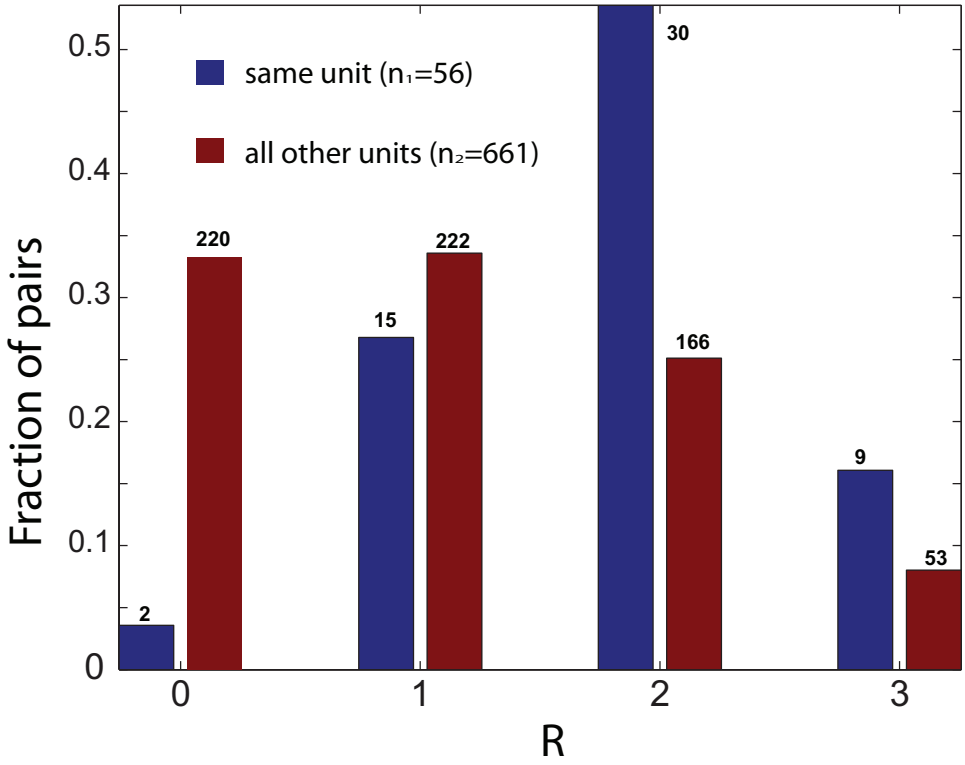

Figure 2.8: Validation using behavioral correlates: Analysis of the behavioral correlate coefficients (R) from 14 session pairs. Plot shows fraction of units for all possible values for R for 56 spike clusters detected as the same (shown in blue), and 661 clusters detected as different (shown in red). Absolute counts for each R value in the two categories is indicated on top of the individual bars. A Wilcoxon ranksum test determined that the two distributions were different ($p=1.05 \times 10^{-8}$).

coefficient on a subsequent session is due to intrinsic changes in the firing rate or variability that cause activity is fall above or below significance.

2.4 Discussion

2.4.1 Accuracy of sorting and consistency of recordings

Neural signatures may change due to intrinsic factors. Hence, in experiments in which there is a significant amount of learning occurring, or different behavioral states (awake vs. anaesthetized, or under some influence of drugs) are being compared it is best to use sessions that are not too widely spaced in time. In addition, there are also extrinsic factors that need to be adjusted or corrected for to ensure that the sessions being tested are comparable. Filter gains and thresholds across sessions must be either held constant or corrected for, as this can affect the final recorded waveshape and amplitude characteristics. Electrode impedances also change over time and these need to be corrected or adjusted for [Nelson et al., 2008] to ensure recordings are consistent. Another important aspect while tracking neurons is that spike clusters should be sorted accurately and consistently across sessions. In our method, we obtain the regression parameters from one session, hence the accuracy and consistency of spike clustering in that session will affect the Type I (false positive) and Type II (false negative) error rate. For example: if a neuron has been split into two separate clusters, it may make the criteria for detection much tighter for the same cluster. This occurs since the logistic regression model may predict the same unit on a subsequent session as a different neuron with a lower probability since it was trained on an incorrectly sorted unit. At the other extreme, when two separate neurons have been combined into one multi-unit cluster, the chances of a Type I error increase, due to the regression being trained on less stringent criteria due to more dispersed metrics.

2.4.2 Other metrics and spike sorting techniques

In this study, we compared only nine metrics to measure neural signatures. Results from single electrode recorded from the motor cortex, and tetrode data recorded from medioprefrontal cortex and amygdala showed that M_1 (Amplitude), M_2 (Shape) and M_5 (ISI shape) were the most salient metrics. Depending on the dataset under question, the choice of metrics can affect performance. For example: data from cells that have very distinct firing characteristics, the spike firing or ISI metric may contribute more to effective separation as compared to spike waveshape or amplitude. In our study, the inclusion of a timing metric (M_5) resulted in better separability. The timing of spikes and its corresponding ISI distribution constitutes an important signature of a neuron which has been used to cluster units [Fee et al., 1996], [Delescluse and Pouzat, 2006]. On the other hand, it has also been shown that ISI distributions are affected by brain state (i.e. awake vs. asleep) [Berke et al., 2004], and thus ISI-based metrics have to be used with caution depending on the data, and especially for data that have substantial state differences between sessions.

Tracking of neurons is the inverse of spike sorting: as opposed to separating clusters within a session based on differences in neural signatures, the attempt in tracking neurons is to combine clusters from two different sessions. Techniques and methods used for spike sorting and clustering can be adapted to perform neural tracking. Some spike sorting techniques actually begin with a much larger set of clusters and then combine them based on similarity [Harris et al., 2000]. Metrics and methods used can be adapted for combining pairs of neurons across sessions. Principal Component Analysis (PCA) is a popular method used for sorting neurons. By using the same principal component vectors for the two discontinuous sessions, cluster overlap can be determined.

2.4.3 Flexibility of the probabilistic approach

The signature of a neuron is not constant due to a number of reasons: recording quality degrades over time due to tissue response; micro-motion of electrodes results in sampling a different population of neurons; there are drifts in the amplitude and shape recorded at a site [Lewicki, 1998],[Snider and Bonds, 1998]. These changes may be due to inherent changes in the neuron, or due to learning or plasticity effects. Neural activity and behavior can be assumed to be reasonably stable on an order of minutes, or up to two days according to a recent study [Chestek et al., 2007]. Our novel method provides a probabilistic estimate that can help to characterize these temporal changes. The modeling method presented accommodates different metrics and the decision boundary can be modified depending on the hypothesis or study in question. For very conservative estimates, the decision boundary can be moved below the default threshold of 0.5. In other cases, for example: while comparing two sessions that are temporally far apart (on the order of weeks) the metrics may be be further away from the origin due to drift, plasticity, or learning-induced changes resulting in Type II errors due to pairs being just above the default-decision threshold of 0.5. In such, cases the decision boundary can be moved and Monte-Carlo cross-validation method can assess the effect of moving the decision boundary. The flexibility afforded by using logistic regression for classification is not possible in the so called 'hard-classifiers', such as linear discriminant analysis (LDA), quadratic discriminant analysis (QDA) or support vector machines (SVMs), that only give the decision boundary but not any associated probability measures [Hastie et al., 2001].

2.4.4 Application of method for data analysis in Chapter III

We applied the method developed in this chapter to obtain an accurate sample size for data analysis in the ChapterIII. For the entire dataset, we tracked spike clusters across sessions of recording and excluded spike clusters that were assessed to be the same or classified as indeterminate. Only spike clusters with a probability of similarity <0.2 were used for the tracking analysis. Refer to section 3.2.6 for details of the analysis.

CHAPTER III

Laminar Analysis of Unit Activity

Abstract

Improving cortical prostheses requires the development of recording neural interfaces that are efficient in terms of providing maximal control information with minimal interface complexity. While the typical approaches have targeted neurons in the motor cortex with multiple penetrating shanks, an alternative approach is to determine an efficient distribution of electrode sites within the layers of the cortex with fewer penetrating shanks. The objective of this study was to compare unit activity in the upper and lower layers of the cortex with respect to movement and direction in order to inform the design of penetrating microelectrodes. Four rats were implanted bilaterally with multi-site single-shank silicon microelectrode arrays in the neck/shoulder region of the motor cortex. We simultaneously recorded unit activity across all layers of the motor cortex while the animal was engaged in a movement direction task. Localization of the electrode array within the different layers of the cortex was determined by histology. We denoted units from Layers 2 and 3 and units as *upper* layer units, and units from Layers 5 and 6 as *lower* layer units. Analysis of unit spiking activity demonstrated that both upper and lower layers encode movement and direction information. Unit responses in either cortical layer of the

cortex were not preferentially associated with contralateral or ipsilateral movement. Aggregate analysis (633 units), tracking analysis (323 units), and best session analysis (75 units) indicated that units in the lower layers (Layers 5,6) are more likely to encode direction information when compared to units in the upper layers (Layers 2,3) ($p<0.05$).These results suggest that electrode sites clustered in the lower layers provide access to more salient control information for cortical neuroprostheses.

3.1 Introduction

In recent years there have been a number of advances in cortical neuroprosthetic devices and methods [Nicolelis, 2001], [Kennedy et al., 2000], [Hochberg et al., 2006], [Velliste et al., 2008]. Several areas in the brain have been shown to encode movement-related signals that potentially could be further developed for neuroprosthetic applications [Scherberger et al., 2005], [Hatsopoulos et al., 1998], [Marzullo et al., 2006], [Taylor et al., 2002]. While candidate neural sources for control signals are being actively explored, the issue of how to best access these signals using implantable microelectrode arrays needs further study. The basic requirement for a broad class of neuroprosthetic devices is the ability to record unit activity reliably for many years and the ability to record action potentials from many different cells within a small volume of cortex [Schwartz, 2004]. Work from various studies concludes that neuroprostheses would benefit from the addition of more channels, with higher recording quality and higher stability [Santhanam et al., 2006], [Patil and Turner, 2008]. However, practical constraints at the level of the neural interface result in cost-benefit tradeoffs. Higher site counts require more complex packaging and multiplexing electronics. Additional penetrating shanks may increase the incidence of deleterious reactive tissue response. A top-down design approach would be to consider interface

designs that are likely to yield the most salient control-related information with the fewest number of electrode sites and penetrating shanks. Anatomically, the neocortex has a regular columnar structure consisting of six layers that is similar across different brain areas [DeFelipe et al., 2002]. Specifically, how does the signal quality of recording in the different layers guide electrode design? How should electrode sites in the six-layered cortex be distributed to maximize salient control information?

Previous neuroprosthetic research has typically targeted neurons in the lower layers of the motor cortex (Layers 5,6) [Serruya et al., 2002], [Taylor et al., 2003], [Donoghue, 2002] to obtain a control signal because the large pyramidal Betz cells in these layers project to the spinal cord, and their large dipole fields result in higher recording quality relative to other cells [Humphrey et al., 1970]. Advances in electrode technology and the ability to conduct long-term, simultaneous, multi-site recordings have made it possible to evaluate event-related action potentials from different cortical layers for movement and direction information. In this study, four rats were implanted bilaterally with single-shank multi-site silicon microelectrode arrays and trained to perform a two-direction movement discrimination task. We used fixed microelectrodes, similar to electrodes used in long-term neuroprostheses, to compare unit activity in the upper layers (2,3) and lower layers (5,6) with respect to movement and direction across and within sessions. To localize the electrode sites in the different layers, we used a combination of selective microlesioning and Nissl-staining of cortical slices for *post mortem* reconstruction of the electrode track. We found units in both the upper and lower layers encode movement and direction, but they do not appear to have an ipsilateral or contralateral direction preference. More importantly, units in lower layers of the cortex are *more likely* to encode directional information as compared to units in the upper layers. Our study suggests that electrodes with

sites clustered in the lower layers will be more effective in obtaining control signals for neuroprosthetic applications.

3.2 Methods

3.2.1 Behavioral Paradigm

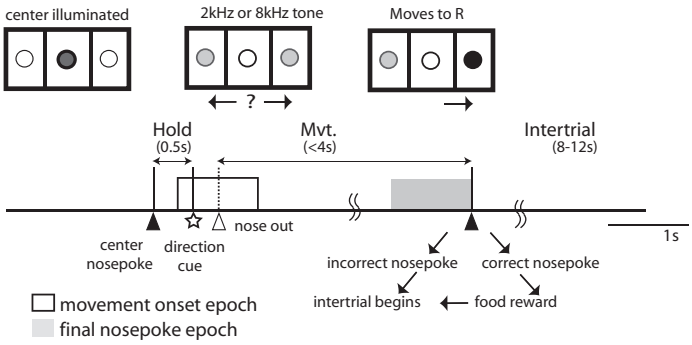

Figure 3.1: Behavioral Paradigm: The behavioral task was a two-direction movement discrimination task. When the center nosepoke was illuminated, the animal self-initiated the task by poking the center nosepoke. After a fixed hold period of 0.5s, a pure tone (2kHz or 8kHz) was played cueing movement to the left or right nosepoke. The animal then inserted its nose into the left or right nosepoke. If the animal failed to hold for the minimum period, the trial was aborted. If the animal correctly moved to the cued nosepoke, it was rewarded with a food pellet. The trial was ended following a correct or incorrect nosepoke. After a variable intertrial period of 8-12s, the center nosepoke was illuminated again to indicate that the next trial could be initiated. The two boxed regions denote the two analysis epochs: the 'movement onset' epoch, which is the 1s window around movement onset shown by the dashed line (variable due to reaction time delay); and the 'final nosepoke' epoch, which is the 1s window before the final nosepoke.

Prior to surgery, animals were food deprived to 85% of their free-feeding weight and trained for 2 to 3 weeks on a two-direction movement discrimination task [Cohen and Nicolelis, 2004] using a three-aperture nosepoke. It should be noted that this is an all-body movement task, which is substantially different from the reaching and grasping tasks typically performed in non-human primate experiments. The behavioral paradigm is shown in Fig. 3.1. At the start of each trial, the center nosepoke was illuminated and the rat self-initiated the trial by inserting its nose into the center nosepoke. Motion into and out of the nosepokes was detected by infra-red

photobeams. After a fixed hold period of 0.5s, a pure tone of 2kHz or 8kHz (in the auditory range of the rat [Otto et al., 2005]) was played in a random sequence which cued the rat to move to the left or right nosepoke, respectively. If the animal failed to hold for the minimum period before the tone, the trial was aborted. If the rat responded by inserting its nose in the cued nosepoke, it was rewarded with a 45mg food pellet (Bio-Serv, Frenchtown, NJ) and the intertrial period began; if the animal responded by inserting its nose in the non-cued nosepoke, the trial was ended and the intertrial period began. After a random intertrial period of 8 to 12s, the center nosepoke was illuminated again to indicate that the next trial could be initiated. Video analysis of the task showed stereotypical movements of nose removal from the center nosepoke and movement of the neck/head to one of the two nosepokes. For trials in which the rats took longer than 4s to respond, the rats tended to pull their nose out and walk around before inserting their nose in the left or right nosepoke. Thus, to avoid confounding the results, all trials with movement times greater than 4s were excluded from analysis. There were 100-150 such trials during each session, and typically one session was run per day. At the end of the session, rats were supplemented with standard rat food pellets to keep them at 85% of their free-feeding weight. Animals were kept on a reversed 12-hour light/dark schedule and run during the dark cycle.

3.2.2 Surgical Implantation and Preparation

Once the behavioral paradigm was mastered by the animal ($>$ 85% correct trials, typically 3 weeks), we implanted the electrode arrays. We did not notice a bias towards one side during the training period for any of our rats. Error rates were similar for both sides. Four male Long-Evans rats (Charles River Labs, Boston, MA) were implanted in the neck region of the motor cortex (M1) with

a single-shank 16-site silicon-substrate microelectrode (c1x16-6mm100-1250, NeuroNexus Technologies, Ann Arbor, MI). All animals had two such implants, one in each hemisphere of the cortex with stereotaxic coordinates: A.P.+3.0, M.L.±2.5 [Donoghue and Wise, 1982]. All arrays were 6mm long, had site sizes of 1250μm^2, and 100μm spacing between each of the sixteen electrode sites. Surgery was performed as previously described in [Vetter et al., 2004] and [Marzullo et al., 2006]. All surgical procedures were carried out with University for Laboratory Animal Medicine (ULAM) and University Committee on Use and Care of Animals (UCUCA) approved protocols at the University of Michigan. Anesthesia was maintained through intraperitoneal injections of a mixture of 50mg/ml ketamine, 5mg/ml xylazine, and 1mg/ml acepromazine at an injection volume of 0.125ml/100g body weight. At every subsequent hour of surgery, 0.1ml ketamine (50mg/ml) was delivered to the animal to maintain anesthesia. Each animal was secured to a standard stereotaxic frame, and four stainless-steel bone screws were inserted into the skull. A stainless-steel ground wire attached to the electrode connector was connected to one bone screw over the cerebellum, as temporary mechanical support until the connector was permanently cemented to the skull using dental acrylic. It also later served as an electrical ground point. A craniotomy (2mm diameter hole) was performed over the target cortical area, and the dura mater was cut away to reveal the cortical surface. The electrode array was inserted by hand using #5 fine PTFE-coated forceps into the target cortical area. Typically, the electrode was inserted such that the top site was just below the cortical surface. The electrode assembly was wrapped with GelFoam (Pfizer, Inc., New York) and then cemented with dental acrylic. The skin around the acrylic was tightened with sutures and anti-bacterial cream was applied. Animals were given 3-4 days to recover post-surgery before experiments were resumed. As a control to

verify if responses were not a result of afferent feedback, we also performed passive left and right nosepoke movements on an additional set of two anaesthetized animals with identical electrode arrays and locations in the motor cortex. Our results indicated that passive driving does not affect the results from this study as there were no significant differences in unit firing rates between left and right movements.

3.2.3 Recording Procedure

After recovery from the surgery, animals continued to perform the same behavioral task while neural activity was recorded from both hemispheres. Spike times, spike waveforms, local field potentials (LFPs), and external events were recorded simultaneously using a Plexon Multichannel Acquisition Processor (MAP, Plexon Inc., Dallas, TX). The signal from each electrode was passed through a high-input impedance headstage with unity gain and then filtered to extract the spike and the LFP components. For spike recordings, the signals were filtered with a passband of 154-8800 Hz, further amplified, and sampled at 40 kHz. Thresholds were manually set and spike waveforms were stored from 150μs before to 700μs after threshold crossing. Behavioral events were sampled at 25μs resolution. Unit activity was sorted offline from each channel using Offline Sorter (Plexon Inc., Dallas, TX). Fig. 3.2 shows sample waveforms, ISIs, and PCA clusters from all animals. Data were recorded in a single session each day, for a period of 4-6 weeks.

3.2.4 Microlesioning and Histology

At the end of the study, we performed microlesioning followed by histological analysis of Nissl-stained cortical slices to determine the electrode site locations within the different layers. First, we measured 1Khz site impedances, and selected three to four sites that had impedances less than 1.5 MΩ and were approximately at

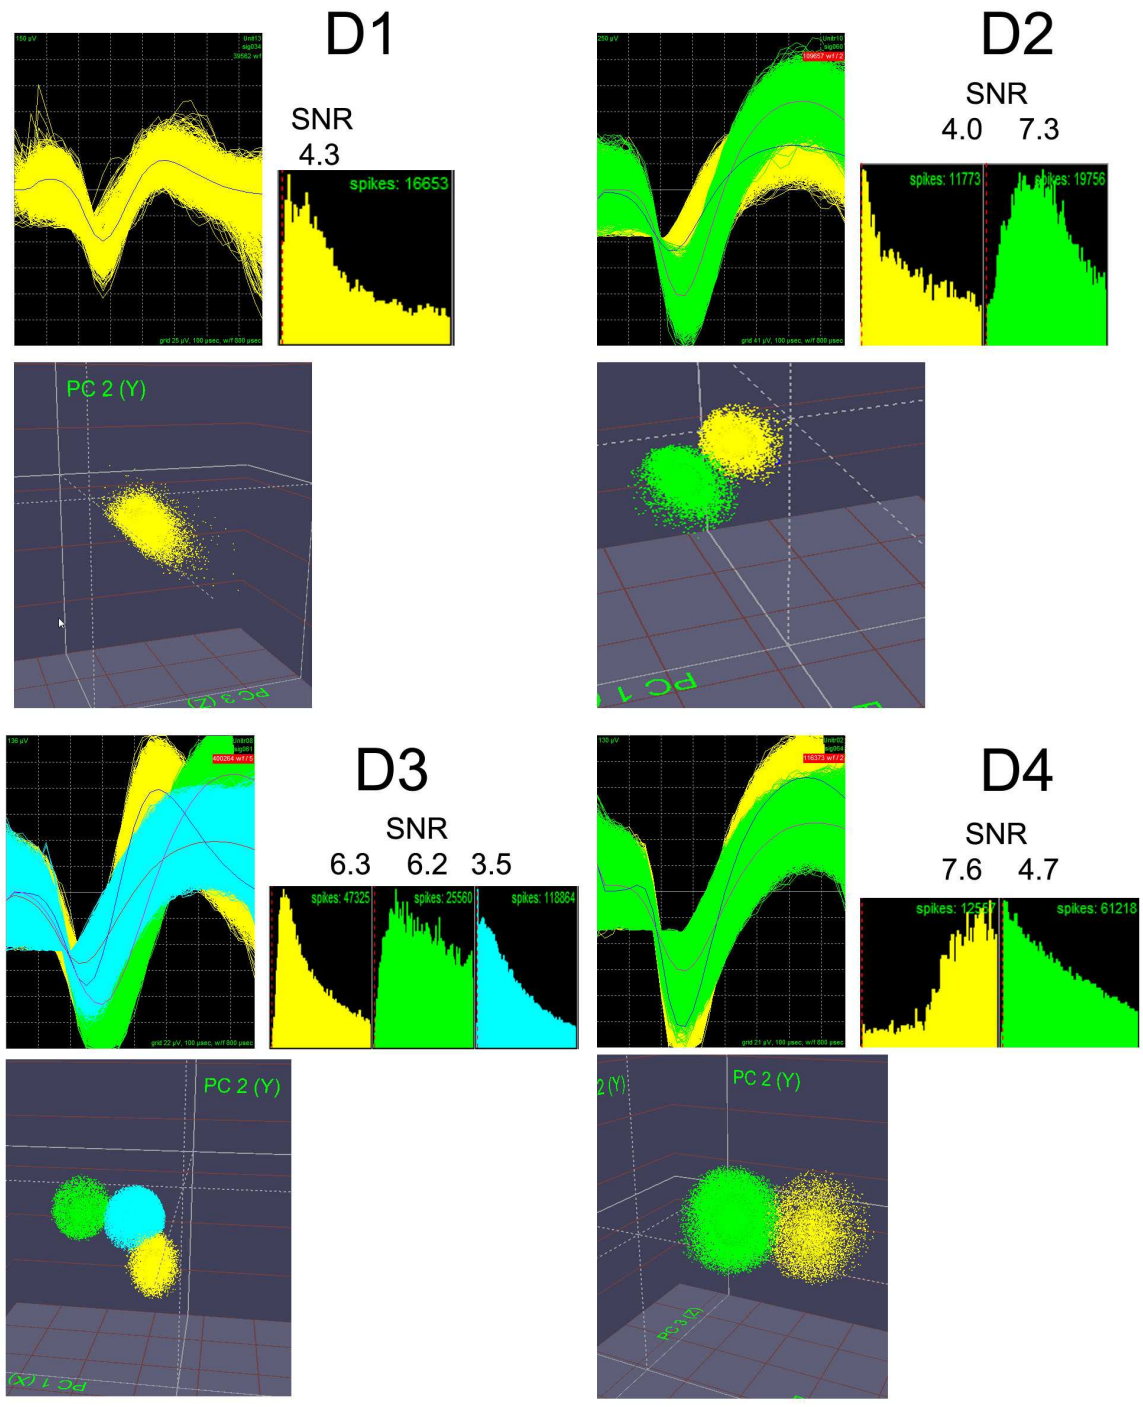

Figure 3.2: Sample Waveforms: Representative waveforms from all four animals D1-D4 showing sorted waveforms, associated signal-to-noise ratios (SNRs), inter-spike intervals (ISIs), and 3-D principal component (PC) clusters of sorted units.

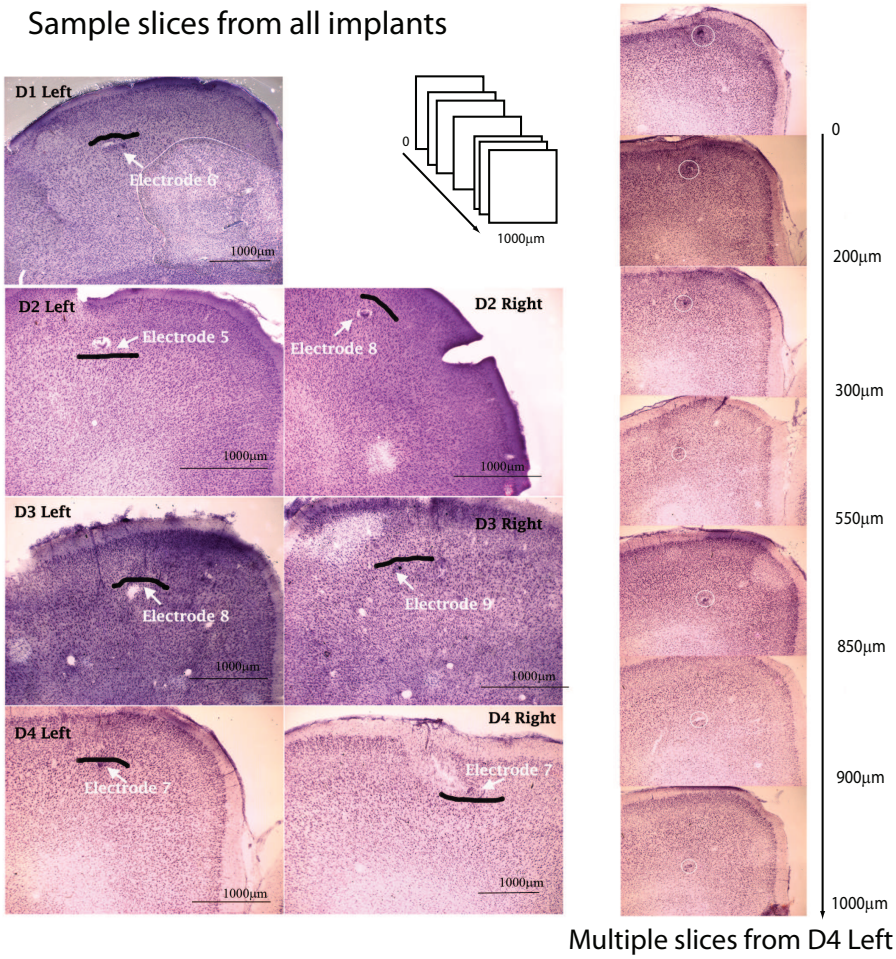

Figure 3.3: Histology and electrode tracks: Left panel shows Nissl-stained coronal sections of sample slices from all animals D1-D4 showing electrode tracks or lesion marks for all seven implantations. The black line marks the boundary between the upper and lower layers. The right panel shows seven coronal sections arranged rostro-caudally, as indicated by the schematic, for one implant (D4 Left) showing alternating lesions and electrode tracks which were used to reconstruct site locations.

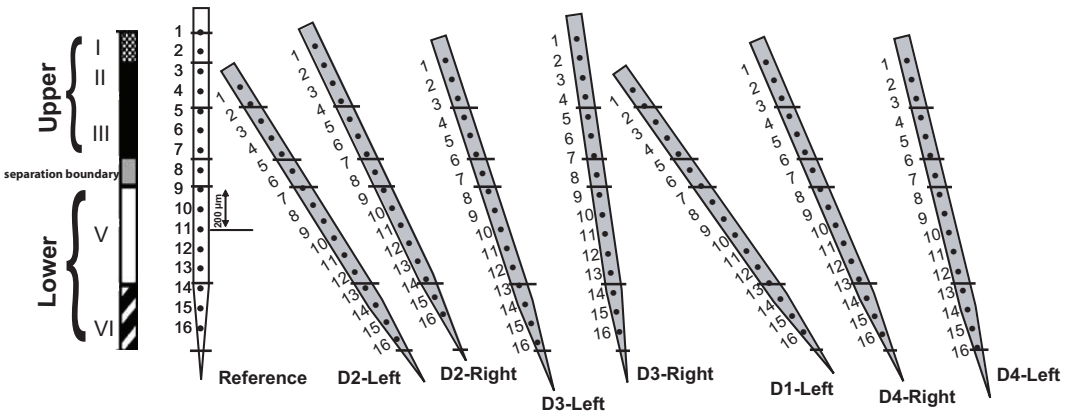

Figure 3.4: Implant Location: Cartoon shows location and orientation of the different electrode sites in the various layers for all animals (layer thicknesses are approximately scaled). The gray band is the 200 μm separation region between the upper and lower layers.

the top, middle, and bottom of each electrode array. At these selected electrode sites we passed 35μA DC for 2s using a potentiostat (AUTOLAB, Eco Chemie, Netherlands) to create micro-lesions (these parameters were determined empirically by unpublished experiments in our lab). We waited 2-3 hours for microlesion 'scars' to form, and then perfused the animal intracardially with 4% paraformaldehyde and explanted the brain tissue. In all cases, electrode arrays extracted from the brain were intact and attached to the skull/headcap. After fixing the tissue, we took 50μm coronal sections and then performed a standard cresyl-violet Nissl stain on the tissue slices. The slices were then analyzed under a microscope and images were taken to reconstruct the position of the electrode array based on the shank track, lesion marks, and the known geometry of the probe. After determining the location of the electrode array, we identified the location of the boundary between the upper and lower layers (the start of Layer 5) by analyzing the gray-level index values of the images using ImageJ (RSB, NIH: http://rsb.info.nih.gov/ij/). At the beginning of Layer 5, the gray-level index increases abruptly due to the large pyramidal cells and the higher cell density.

For the purpose of this study, we denote units collected from Layers 2 and 3 as *upper* layers, and units from Layers 5 and 6 as *lower* layers. Layer 1 contains few recordable cell bodies, and Layer 4 is small and its existence in the rat motor cortex is debated [Brecht et al., 2004]. Data from the site that was closest to the determined boundary between the upper and lower layers was not used for subsequent analysis (this effectively created a 200μm separation between the upper and lower layers). This was done to account for the margin of error in determining the precise boundary between layers, leading to a possible incorrect assignment of a site at the boundary to a different layer. Fig. 3.3 shows array tracks, electrode lesion marks, and the boundary between the upper and lower cortical layers in brain slices from all animals (D1-D4), and multiple sections for one implant (D4 Left). The D1 Right implant is not shown since that array was found to be non-functional after implantation.

3.2.5 Sorting quality based on signal-to-noise ratio (SNR)

Units were manually sorted based on a number of different criteria: principal component cluster analysis, auto- and cross-correlograms, inter-spike intervals (ISIs), and wave-shape. We calculated the signal-to-noise ratio (SNR) based on the individual waveforms of the different units on a channel as an additional quantitative metric. SNR was based on the individual waveforms given by the formula:

$$SNR = \frac{A}{2 * SD_{Noise}} \tag{3.1}$$

where A is the amplitude of the peak-to-peak voltage of the mean waveform, and SD_{noise} is the standard deviation of the noise (See [Suner et al., 2005] for details of the method). As per the 4-point scale by Suner, et al., units with SNRs between 2 and 4 are considered to be 'fair'; units with SNRs $\geq$ 4 are deemed to be 'good'. Units with SNRs below 2 were classified as 'poor' quality and were not considered

for subsequent analysis.

3.2.6 Measure of task-related significance

To compare unit activity relative to the task, we computed the firing rates of the units in two analysis epochs – movement onset epoch (1s around movement onset) and final nosepoke epoch (1s before the final correct nosepoke and ensuing delivery of food). In the movement onset epoch, we determined units that showed a statistically significant modulation in their firing rates by comparing the 500ms pre-movement and 500ms post-movement firing rates in each trial. A Mann-Whitney U-test (or Wilcoxon rank sum test) was performed to determine if there was a significant difference between the firing rates under the two different conditions for each unit (using the procedure outlined in [Samejima et al., 2005]). Similarly, to determine directional information we compared the unit firing rate differences between ipsilateral and contralateral movements in the entire 1s movement onset epoch and 1s final nosepoke epoch. Only correct trials were used for analysis. We also verified results using peri-stimulus time histograms (PSTHs). Spikes in the PSTHs were smoothed using a Gaussian kernel of 50ms and 95% confidence intervals for the error bars were calculated using a bootstrapping procedure based on the `psth` routine (Chronux Project, http://www.chronux.org). While combining units from a layer for an aggregate PSTH, the individual PSTHs were normalized by the peak firing rate in the epoch under analysis. Across the 2-3 weeks of recording, unit activity at each electrode varied from day to day, presumably due to probe micromotion or due to changes in local environment as a result of probe implantation [Subbaroyan et al., 2005]. For aggregate analysis, we consider units from separate sessions as different units. However, this assumption is not strictly true as the same units could be present across multiple sessions. We used the method developed in Chapter II to obtain an accu-

rate sample size by tracking neurons in successive sessions and excluding units that had a probability of similarity >0.2. We also additionally performed 'best' session analysis that only considers a single session from each electrode. For each channel, the session with the highest SNR for a unit on that channel was chosen as the 'best' session. In this analysis there are no overlapping units, as only one session per site was considered.

3.2.7 Significance of difference between Upper and Lower Layers

After determining which units showed task-related significance, we prepared a 2x2 contingency table layer (upper or lower) versus response (modulated, did not modulate). We sorted and summed all the units in the four categories. This analysis reveals if there is a statistically significant effect of the layer location on the response characteristics of the units. For the aggregate and tracking analysis, we used a paired chi-squared (χ^2 test) to determine if responses from the two populations were statistically significant. For the best session analysis, we used Fisher's exact test to assess significance since the sample sizes were small. All reported p-values are two-tailed values.

3.2.8 Analysis of Direction Preference

We further investigated whether units in either of the two layers had a preference for movement in the contralateral or ipsilateral direction. Each unit analyzed either encoded: no direction, contralateral direction, or ipsilateral direction. We performed multinomial logistic regression using the different direction preferences as categorical variables and denoted the upper and lower layers as regressors (0 and 1 respectively) to predict the probability or odds ratio of observing a particular direction. A significant regressor would indicate an ipsilateral or contralateral direction preference

depending on the layer from which the neuron was recorded. We calculated the regression coefficients, odds ratios, and p-values (`mnrfit` routine, MATLAB, Mathworks Inc., Natick, MA). Sample odds ratio was simplified, and cast in terms of the log probability of observing a contralateral preference as shown below:

$$\ln[\text{Odds ratio}] = \beta_0 + \beta_1(\text{layer}), \tag{3.2}$$

$$\ln\left(\frac{p(\text{contralateral})}{1-p(\text{contralateral})}\right) = \beta_0 + \beta_1(\text{layer}) \tag{3.3}$$

$$p(\text{contralateral}) = \frac{e^{\beta_0+\beta_1(\text{layer})}}{1+e^{\beta_0+\beta_1(\text{layer})}} \tag{3.4}$$

Here, the odds ratio predicts the relative probability or relative odds of observing a contralateral preference with respect to an ipsilateral preference given that these observations were made in the upper or lower layer respectively. This odds ratio is predicted by the regressors β_0 (effect independent of layer) and β_1 (contribution to effect of the layer). Thus, for significant contribution of a layer for contralateral preference, the absolute value of β_1 should be large and significant, and the odds ratio positive. If β_0 is large and significant in comparison, we can conclude that most of the contribution to the side preference is independent of the layer under consideration.

3.3 Results

3.3.1 Site Locations

Fig. 3.4 depicts a schematic of the site locations of all the electrodes in the different layers. Layer thicknesses in the schematic are as per [Hutsler et al., 2005] and are approximately to scale. One array in animal D1 in the right hemisphere was found to be non-functional. All following analyses are based on seven implantations

in the animals D1-D4. Site locations were found to be well distributed across all layers.

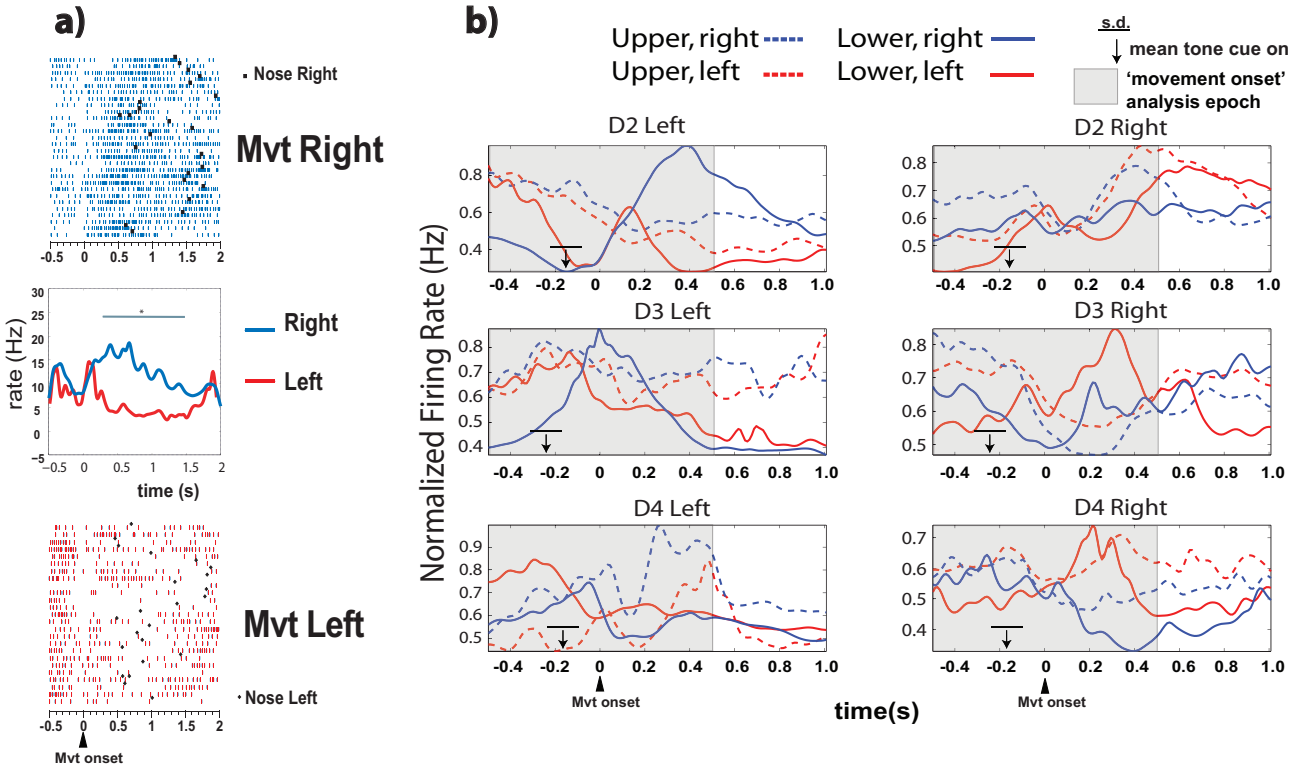

Figure 3.5: (a) Top and bottom panels: Raster plots of a typical unit from implant D2 Left for all trials separated by movement to the right and left respectively. Dots indicate the time of the final nosepoke. Middle panel: Event-triggered PSTH. Bar denotes time period where there was significant difference in firing rate between the two conditions (corrected for multiple comparisons). (b) Normalized PSTHs for units that encoded contralateral movement in the movement onset epoch. PSTHs were normalized by the maximum firing rate. Trials were aligned to the start of movement indicated by the black triangle at t=0; the movement onset epoch analysis window is shown in gray. The tone cues were distributed around the mean offset indicated by the arrow, bar denotes the standard deviation.

3.3.2 SNR Quantification

'Aggregate' analysis pooled units from all channels and sessions. All units used for analysis had SNRs at least ≥ 2.0 and no other criteria was used to exclude units recorded from analysis. For the upper layers, the SNR was **mean=4.51, sd=1.77** and for the lower layers the SNR was **mean=4.42, sd=1.75**. The best session was

determined to be the session which had the largest SNR for a particular channel across all sessions for a particular electrode. Mean SNR was **5.72** and **5.95** for upper and lower layers respectively. There was no statistical difference in the quality of units from the two layers in both best session and aggregate analysis. Thus, there were no recording quality differences between units recorded from the upper and lower layers that would affect subsequent results.

3.3.3 Movement and Direction Related Activity

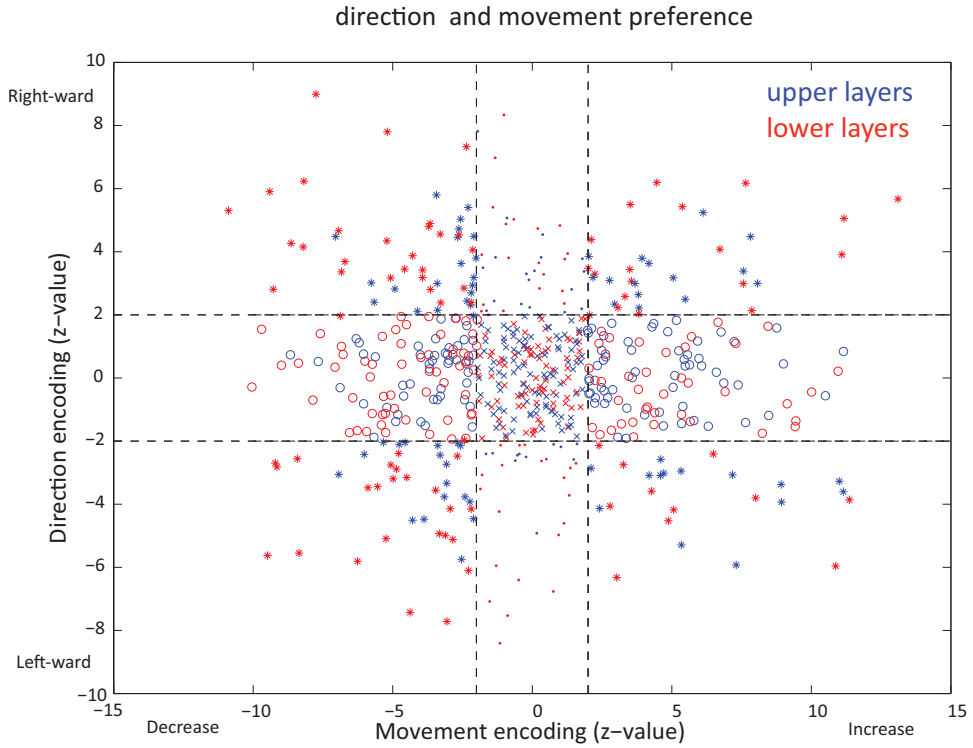

Figure 3.6: Scatter plot shows the z-values for differential firing rates in the analysis windows for movement encoding (pre vs. post) on the X-axis, and direction encoding in the 'movement onset' epoch on the Y-axis for the entire aggregate analysis dataset. Units in the upper layers are shown in blue and units in the lower layers are shown in red. Crosses indicate units that encoded neither movement nor direction. Open circles denote units that showed movement encoding. Dots indicate unit that encoded direction, but not movement. Asterisks indicate units that encoded both movement and direction.

Only correct trials with a response time less than 4s post-tone were used for analysis. Overall behavioral performance across all the sessions used for analysis was as follows — D1: 85.4% correct, D2: 96.5% correct, D3: 75.7% correct, D4:

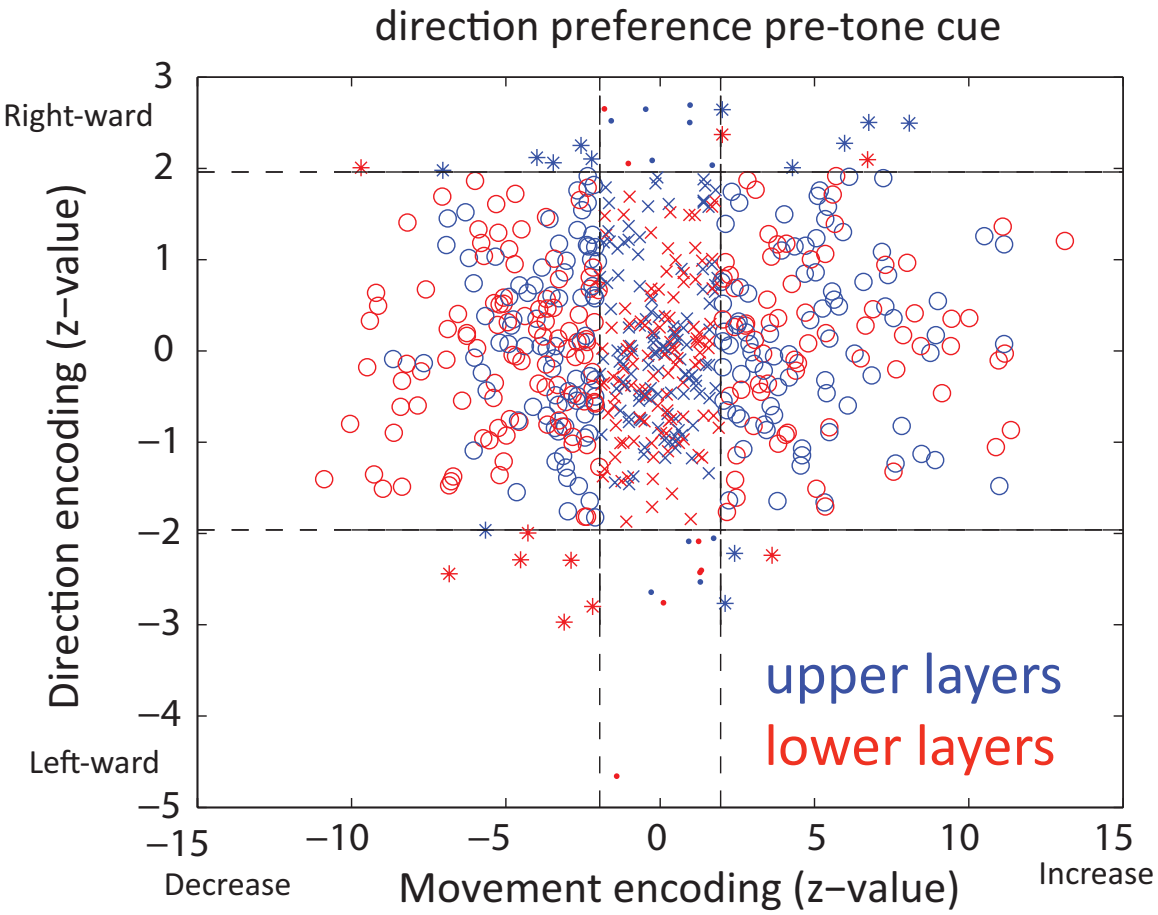

Figure 3.7: Legend and symbols as in the previous plot. Analysis of direction encoding was performed in the 500ms before tone cue for the entire dataset. Movement analysis was performed as in the previous plot. As expected, most units do not encoded direction before the tone cue.

86.2% correct. We investigated the performance of the layers across sessions for the fixed electrodes in the *aggregate*, *tracking* and *best session* analysis. The aggregate analysis considered all units recorded from all channels and sessions as independent units; we recorded from 313 units in the upper layers and 320 units in the lower layers from the seven implantations. In the tracking analysis, all units that had a probability of similarity >0.2 were excluded from the analysis. We analyzed 188 units from the upper layers and 205 units from the lower layers. We also performed a *best session* analysis that considered units within a single session for that channel. For the best session analysis, we had 33 units from the upper layers and 42 units from the lower layers. For both sets, we analyzed unit activity in two different epochs: the movement onset epoch (1s period around the onset of movement) and before final nosepoke (1s period before food delivery) as shown in Fig. 3.1 to determine if encoding in the two layers changed as a consequence of behavioral state. Fig. 3.5a shows sample raster plots and PSTHs for one sample unit. Fig. 3.5b shows PSTHs for all animals; drop arrows and the bars correspond to the mean reaction time and associated standard deviation respectively. We followed the non-parametric permutation approach for significance testing [Womelsdorf et al., 2006] to correct for multiple comparisons in Fig. 3.5 a. Figure 3.6 shows the z-values of differential firing rates in the $\pm$500ms analysis window for movement encoding (pre vs. post) and direction encoding (movement onset epoch) for units in the upper layers versus lower layer. Figure 3.7 shows z-values for whole dataset when the direction encoding analysis in left and right trials was performed in the 500ms epoch before the onset of the tone cue. Analysis of movement encoding is on shown on the X-axis, and was carried out as described earlier. As expected, a majority of units did not encode direction before the onset of the cue. The small fraction of units that seem to encode

direction before the onset of the tone cue may be due to chance.

Movement encoding:

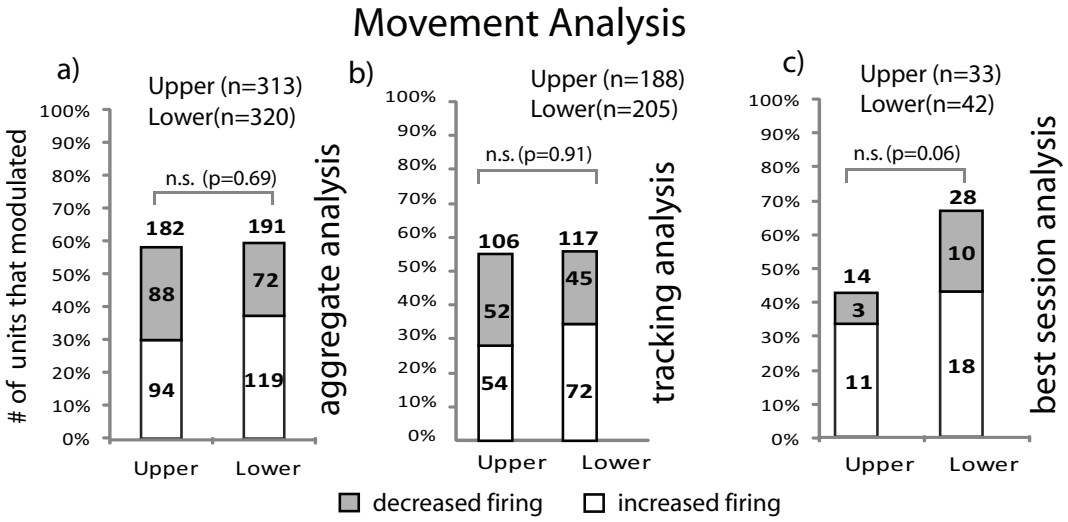

Figure 3.8: Movement Encoding: Bar graphs show the percentage of units that showed an increase or decrease in firing rate with respect to layer. (a) Aggregate analysis shows no statistically significant difference (p=0.69) in the total number of units that modulated activity between the upper (n=313) and lower layers (n=320). Upon consideration of the kind of modulation, units in the lower layers were significantly more likely (p=0.04) to increase than decrease their firing rate. We were unable to detect any such preference for the modulating units in the upper layers. (b) Tracking analysis also shows no statistically significant difference (p=0.91) in the total number of units that modulated activity between the upper (n=188) and lower layers (n=205). (c) Best session analysis shows no statistically significant difference (p=0.06) in the modulation of unit firing rate between the upper (n=33) and lower layers (n=42) with respect to movement.

When we compared firing rates of units in the movement onset epoch, both upper and lower layers contained units that increased or decreased firing with respect to movement and direction. Fig. 3.8 summarizes results from the aggregate, tracking and best session analysis with respect to movement encoding. We determined that there was no significant difference between units in the upper and lower layers with respect to movement in the aggregate analysis ($p = 0.69$) and tracking analysis ($p = 0.91$). Units in the lower layers had a significantly greater tendency to increase firing rate with respect to movement ($p = 0.04$). We observed that 119 out of 191

(62.3%) units in the lower layers showed an increase in firing as compared to 94 out of 182 (51.6%) units in the upper layers in the aggregate analysis. In the best session analysis we found that the difference between upper and lower layers in terms of encoding movement was also not statistically significant (p = 0.06).

Direction encoding:

Fig. 3.9 summarizes results from aggregate and best session analysis with respect to direction encoding for both the 'movement onset' and 'final nosepoke' epoch. When we compared direction encoding in the aggregate analysis, the lower layer units were found to be significantly more likely to encode direction as compared to units in the upper layers in both the movement onset (p=0.03) and final nosepoke epoch (p=0.0002). In the tracking analysis, units in the lower layers were found to be significantly more likely to encode direction of the movement as compared to units in the upper layers in the movement onset (p=0.01) and final nosepoke epoch (p=0.003). In the best session analysis, units in the lower layers were found to be significantly more likely to encode direction of the movement as compared to units in the upper layers in the movement onset (p=0.03) and final nosepoke epoch (p=0.02).

It has been reported that some neurons in the motor cortex respond to non-spatial sensory cues and are independent of movement direction [Salinas and Romo, 1998]. We compared neural activity during correct (hit) and incorrect (error) trials using a permutation test, as described in [Salinas and Romo, 1998] and [Britten et al., 1996], on all the directionally-tuned neurons in sessions with at least five error trials in the two error classes. We were not able to perform this analysis on any of the sessions from rat D2 since the error rates were low (incorrect trials $< 4\%$). Upon analysis of the data, we found that 4 out of 38 (10.5%) neurons from the upper layers, and 4 out of 41 (9.7%) in the lower layers were found to be encoding sensory rather than motor

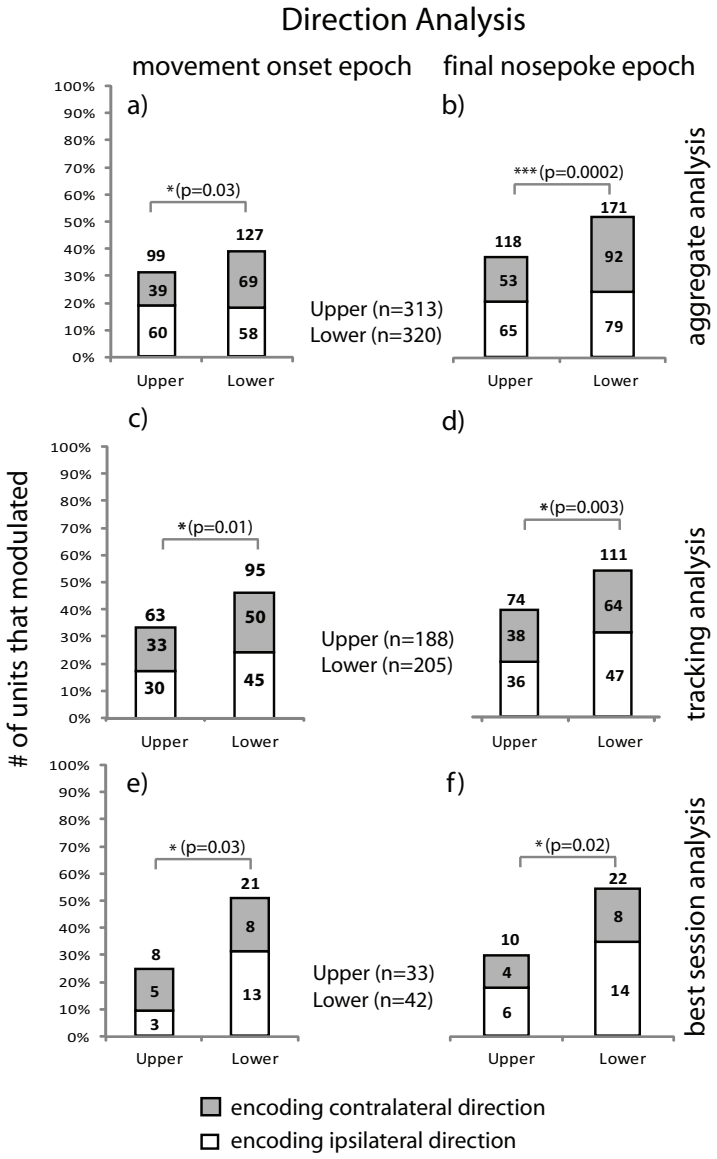

Figure 3.9: Direction Encoding: Bar graphs show the percentage of units that showed an ipsilateral or contralateral direction preference with respect to layer. (a) and (b) Aggregate analysis shows statistically significant differences in the movement onset (p=0.03) and final nosepoke epochs (p=0.0002), in modulation of unit firing rate in the upper (n=313) and lower layers (n=320) with respect to direction encoding. (c) and (d) Tracking analysis shows statistically significant differences in the movement onset (p=0.01) and final nosepoke epochs (p=0.003), in modulation of unit firing rate in the upper (n=188) and lower layers (n=205) with respect to direction encoding. (e) and (f) Best session analysis shows statistically significant differences in the movement onset (p=0.03) and final nosepoke epochs (p=0.02), in modulation of unit firing rate in the upper (n=33) and lower layers (n=42) with respect to direction encoding.

information. After accounting for these neurons, the conclusions remain unchanged.

Ipsilateral versus contralateral encoding:

We further investigated whether units in the upper or lower layers had a preference for ipsilateral versus contralateral movement. We hypothesized that the lower layers would show a contralateral preference given the large number of output neurons to the contralateral side in the lower layers of the cortex. Among the direction encoding units, we found that both layers encode ipsilateral and contralateral movements. For the log odds ratio of contralateral versus ipsilateral direction given by Equation 3.2, we obtained $\beta_0 = -0.32$ (p=0.09) and $\beta_1 = 0.32$ (p=0.20) for the movement onset epoch; and $\beta_0 = -0.20$ (p=0.27) and $\beta_1 = 0.32$ (p=0.13) for the final nosepoke epoch. We found units in both the upper and lower layers encode movement and direction, but do not appear to have a specific ipsilateral or contralateral direction preference. Data from this study suggests that both ipsilateral and contralateral direction can be obtained from a single cortical implant in one hemisphere. This lack of a contralateral bias could be due to the implants being in the neck/head region which are more medially located and hence can be expected to have more of a bilateral drive.

3.4 Discussion

3.4.1 Implications for electrode and algorithm design for neuroprosthetics

A viable long-term neuroprosthetic device requires the recording electrode arrays to be chronically implanted and held fixed. The objective of the aggregate analysis was to assess the performance of chronically implanted penetrating microelectrodes across sessions. Apart from requiring that units have an SNR above 2.0, no other selection criteria was applied to the units recorded across the different sessions. Thus,

the aggregate analysis presents pooled results from different sessions under the realistic constraint that electrodes, once implanted, are not adjusted. Since the electrodes were held fixed, units across sessions for a given implant are not necessarily independent. Therefore we also investigated movement and direction performance of units in the two layers considering only a single session for each site. Best and aggregate analysis showed that units in both upper and lower layers encode movement. About 60% of the units recorded from both layers encoded discernible movement information, and there was no statistically significant difference between them. Upper layers showed about equal number of units that increased and decreased firing with respect to movement; in contrast, units in the lower layers had a tendency to increase firing. This difference could be due to neurons in each layer being involved in different cortical microcircuits.

A simple model of cortical processing hypothesizes that the superficial pyramidal neurons combine feedforward input from subcortical, inter-areal, and intra-areal excitatory sources and explore possible interpretations, whereas the deeper layers are organized to exploit the evolving interpretations of these signals [Douglas and Martin, 2004]. The layer 5 pyramidal cells process the local superficial signals and converge their output to motor structures. The outputs also drive subcortical structures such as the basal ganglia, red nucleus, cerebellum, and ventral spinal cord. Recent research suggests this descending chain of events by demonstrating that the upper layer neurons act like preamplifiers driving output neurons in the lower layers [Weiler et al., 2008]. This implies that the upper layer neurons play the role of network controllers as they inhibit and excite downstream neurons, while lower layer neurons select specific motor outputs. Our results suggest that lower layers are more salient targets for control signals as this will maximize control

information per electrode site. In terms of single shank electrodes, we suggest that sites should be clustered at the end of the shank to record from deeper layers instead of being uniformly spread along the shank. This does not necessarily require an increase in length of the probe, but an alteration of site spacing.

Several studies have found a variety of arm directions encoded by neurons recorded with a single multi-site implant in one hemisphere [Taylor et al., 2002], [Velliste et al., 2008]. In our study, when the direction preference was considered, both layers showed similar ipsilateral and contralateral tuning. Since predominant motor cortical output projects to the contralateral side of the spinal cord, we expected there would be a similar preference for the contralateral side. We found units in both the upper and lower layers do not appear to have an ipsilateral or contralateral direction preference. This is encouraging from a neuroprosthetic point of view as it precludes the necessity of implanting microelectrode arrays in both hemispheres to obtain effective ipsilateral and contralateral control, although this study only investigated a two-directional task. In our study, the implants were in the head/neck region and hence can be expected to have more of a bilateral drive which may explain the absence of an observed contralateral bias. We did not consistently observe ipsilateral cells decreasing firing or contralateral cells increasing activity. It has been previously observed in non-human primates that for unimanual movements, ipsilateral cells have a tuning response similar to that for contralateral movements; for bimanual movements, ipsilateral cells have preferred directions randomly shifted [Rokni et al., 2003]. Thus, there may not necessarily be a decrease in firing for ipsilateral movements. This study used a two-directional whole-body movement task in the rat motor cortex. Additional experiments would need to be performed in a multidirectional, 3-D task using non-human primates to verify these results and

parse additional movement parameter information such as velocity, acceleration, joint dynamics, etc. for human neuroprosthetic use.

3.4.2 Minicolumns and functional extent

From the point of view of a neuroprosthetic system, accurate characterization of architecture of the neocortex will enable better design of probe geometries to increase throughput of signals. The six-layered structure of the neocortex is roughly similar across brain areas; only the relative thickness of the layers, number of neurons, and cell type differs [DeFelipe et al., 2002]. Observations in the cat somatosensory cortex led Mountcastle to hypothesize the concept of a minicolumn [Mountcastle et al., 1957] which extends perpendicularly across all six layers and forms a basic functional unit with similar response properties. But, properties of cells have found to vary across a single orientation column in the visual cortex [Bauer et al., 1983]. While histological analysis revealed that all our implantations were at an angle, the optimal angle of implantation for a penetrable microelectrode in the motor cortex is yet to be determined.

In the motor cortex, the distribution of directionally tuned cells is non-uniform and highly structured in both dimensions. Specifically, M1 cells with similar preferred directions tend to condense in a vertical dimension forming an ordered structure of minicolumns perpendicular to the surface of the neocortex of width $\approx 30\mu$m and repeating at a lateral distance of 200μm [Amirikian and Georgopoulos, 2003], [Georgopoulos et al., 2007]. According to the Amirikian et al. study, in the monkey motor cortex the functional extent depthwise was 500μm (dorsal-ventral) compared to the cortical column length of 2200μm; whereas in our study, we observed that the functional extent spans almost the entire vertical length of rat motor cortex as units from Layers 2 through 6 are involved in encoding movement and direction. This

difference in functional extent could be due to inter-species difference, since the number of neurons contained in a vertical cylinder of cortical tissue varies across species [DeFelipe et al., 2002]. In the Amirikian et al. study, the sites were not well localized in the vertical direction and could have contributed to the effect of seeing a smaller resultant functional distance. Recent work by the same group has demonstrated a novel method to determine electrode locations using fluorescent dyes and registering Nissl-stained slices to investigate the organization of preferred directions in the motor cortex [Naselaris et al., 2005], [Naselaris et al., 2006]. Using this technique they recently suggested that two kinds of pyramidal cells exist in M1, and interneurons dynamically alter the preferred directions of one class of pyramidal cells which are present across layers of the cortex affecting directional information processing during the preparation and execution of reaching movements [Merchant et al., 2008]. These advances in techniques to determine the location and type of cell recorded in an awake, behaving preparation via extracellular recordings will help validate proposed cortical microcircuits and the functional role of the different cells across the various layers [Du et al., 2008].

3.5 Conclusion

Units in both the upper and lower layers of the rat motor cortex encode movement and direction information. Analysis across sessions and within sessions showed that units in the lower layers are significantly *more likely* to encode direction information as compared to units in the upper layers. These results suggest that electrode geometries with sites clustered in the lower layers will provide access to more salient control information. Improved cell and layer labeling techniques, and more degrees of freedom in a behavioral task will enable us to further parse details of movement and

direction encoding in the motor cortex. These results encourage further investigation into utilizing layer-specific differences in the context of a human neuroprosthesis.

CHAPTER IV

Laminar Analysis of LFP Activity

Abstract

Local field potentials (LFPs) have been shown to encode movement, direction, preparation and attentional features, and therefore can be used as control signals for neuroprosthetic control. From the standpoint of long-term neuroprostheses, they are a suitable alternative to unit activity as they less affected by glial encapsulation, and hence are relatively more stable. In this study, we investigate LFP activity to determine if there are significant interactions in time and/or frequency across the different layers that can be used to obtain more effective control signals. Four rats were implanted bilaterally with multi-site, single-shank silicon microelectrode arrays in the neck/shoulder region of motor cortex of three rats. We simultaneously recorded LFP activity across all layers of the motor cortex while the animal was engaged in a movement-direction task. Localization of the electrode array within the different layers of the cortex was determined by histology. We performed both evoked potential and spectral analysis of LFPs in four frequency ranges – low (3-15Hz), beta-gamma (15-40Hz), high-gamma (40-70Hz), and high (>70Hz) – across both upper (Layers 2,3) and lower layers (Layers 5,6) of the cortex. Our analysis based on 585 LFP recordings from 39 sessions shows that the low frequency range (3-

15Hz) contains more directional information as compared to other frequency ranges. We also found a significant difference in LFP activity between the upper and lower layers of cortex in the high-gamma (40-70Hz) range, but not in the other frequency ranges. Our results indicate that LFPs are viable alternative control signals that can be recorded from either upper or lower layers of the cortex for performance comparable to our results from unit activity.

4.1 Introduction

While advances in microelectrode technologies and neural interfaces have greatly improved the performance and capabilities of single electrode recordings [Seymour and Kipke, 2007], [Kipke et al., 2008], it is difficult to maintain single unit recordings consistently and reliably. For a practical, functional human neuroprosthesis, demonstration of reliability and longevity of control signals over a span of months to years is essential [Donoghue, 2002], [Schwartz, 2004]. Also, for greater degrees of freedom, future systems will need to rely on a number of different brain signals and maximize the theoretical control information that can be obtained [Andersen et al., 2004]. Local field potentials (LFPs) have emerged as suitable alternative control signals as they have been shown to encode features of movement, direction, preparation and attention in various different regions of the brain [Scherberger et al., 2005],[Rickert et al., 2005], [O'Leary and Hatsopoulos, 2006]. As alternative control signals LFPs possess a few advantages: they are less prone to effects of glial encapsulation and electrode micromotion, making them relatively more stable than unit activity; they are hypothesized to possess information independent from the individual spiking activity, as they represent a summed potential from a number of neurons; they do not require any

additional neural interface as they can be simultaneously recorded from the same electrodes used to record unit activity.

While the exact origin of LFPs is unknown, they are commonly understood to be complex resultants of underlying synaptic activity, afterpotentials of somadendritic spikes, and activity of interneurons [Logothetis et al., 2007]. For the same reason, it is believed that LFPs may not reflect the transient and specific synchronization of smaller groups of neurons, and thus lack the spatial and temporal resolution of unit recordings. LFPs may consequently play a more global role in the context of motor planning or preparatory function, while the details of motor action are encoded in the unit activity [Donoghue et al., 1998]. Evoked potential analyses of LFPs have revealed significant LFP activity in response to sensory or motor events [Donchin et al., 2001], [Mehring et al., 2003]. Different frequency bands of LFPs have been shown to encode various parameters of movement. Oscillations in the gamma range (25-90Hz) were shown to modulate with saccadic eye movements in the parietal cortex [Pesaran et al., 2002]. Fast oscillations of motor cortical LFP activity have been investigated previously in a number of studies in monkeys. During exploratory behavior, 25-35Hz oscillations were reported [Murthy and Fetz, 1992]; 15-50 Hz oscillations were found to increase in relation to movement preparation and decrease during movement execution [Donoghue et al., 1998]; 20-30 Hz oscillatory activity of LFPs was reported to occur increasingly during maintenance of a precise grip [Baker et al., 1999]. It was recently demonstrated that movement evoked potentials (MEPs) in the <4Hz, 6-13Hz and 63-200Hz range had the best decoding ability, while the 16-42Hz range provided little directional information [Rickert et al., 2005]. In another study, phase locking and directional tuning were observed in the low (<10Hz) frequency band, and weak decoding was observed in the fast (25-40Hz)

and intermediate (10-25Hz) frequency bands [O'Leary and Hatsopoulos, 2006]. Preferred directions of LFPs have a non-uniform distribution and different frequency bands showed inconsistent direction preferences [Asher et al., 2007]. Thus, analyses have shown significant task related-LFP activity across the entire spectrum from 0.1Hz to 200Hz.

While spectral and temporal characteristics of LFPs have been investigated, these observations have not been studied adequately in conjunction with electrode location within the layers of the cortex. Some of the differences observed in previous experiments may be due to differences in tasks and behavioral paradigms. The observed differences in LFP characteristics could also be due to dissimilar locations of the recording electrodes within the layers of the cortex across experiments. Proposed cortical microcircuit models have hypothesized that the upper layers (Layers 2,3) perform feature selection after integrating sub-cortical and inter-areal input, while the lower layers (Layers 5,6) are involved in actual output [Douglas and Martin, 2004]. Stimuli targeting the upper, but not lower layers, evoked network wide events implying that upper layer neurons drive output neurons in lower layers [Weiler et al., 2008]. A recent study showed that LFPs are local phenomena occurring due to synaptic or neural activity within 250μm of the recording electrode [Katzner et al., 2009], and not on the order of millimeters as suggested earlier by [Mitzdorf, 1985]. Based on the above, we investigated differences in evoked potential and spectral properties of LFPs in the upper and lower layers of the cortex. We hypothesize that LFP activity in the upper layers is more indicative of integration of input and thus will provide more directional information than LFPs in the lower layers. Using single-shank, multisite electrodes we recorded LFP activity across all layers while animals were engaged in a movement-direction task. Localization of the electrode array within the different

layers of the cortex was determined by histology.

4.2 Methods

4.2.1 Behavioral Paradigm and Surgical Procedure

Details of the behavioral paradigm and surgical procedure are described in Section 3.2 and the behavioral paradigm is as shown in Figure 3.1. Animals in this study are same as the ones used in the previous study. Briefly, animals were trained for 2 to 3 weeks on a two-direction movement discrimination task [Cohen and Nicolelis, 2004] using a three-aperture nosepoke. At the start of each trial, the center nosepoke was illuminated and the rat self-initiated the trial by inserting its nose into the center nosepoke. After a fixed hold period of 0.5s, a pure tone of 2kHz or 8kHz (in the auditory range of the rat [Otto et al., 2005]) was played which cued the rat to move to the left or right nosepoke, respectively. If the animal failed to hold for the minimum period before the tone, the trial was aborted. If the rat responded by inserting its nose in the cued nosepoke, it was rewarded with a food pellet and the intertrial period began; if the animal responded by inserting its nose in the non-cued nosepoke, the trial was ended and the intertrial period began. After a random intertrial period of 8 to 12s, the center nosepoke was illuminated again to indicate that the next trial could be initiated. There were 100-150 trials during each session, and typically one session was run per day. At the end of the session, rats were supplemented with standard rat food pellets to keep them at 85% of their free-feeding weight. Animals were kept on a reversed 12-hour light/dark schedule and run during the dark cycle.

Once the behavioral paradigm was mastered by the animal ($> 85\%$ correct trials, typically 3 weeks), we implanted the electrode arrays. Three male Long-Evans rats (Charles River Labs, Boston, MA) were implanted in the neck region of the motor cortex (M1) with a single-shank 16-site silicon-substrate microelectrode (c1x16-

6mm100-1250, NeuroNexus Technologies, Ann Arbor, MI). All three animals (D2-D4) had two such implants, one in each hemisphere of the cortex with stereotaxic coordinates: A.P.+3.0, M.L.±2.5 [Donoghue and Wise, 1982]. All arrays were 6mm long, had site sizes of 1250μm^2, and 100μm spacing between each of the sixteen electrode sites. All surgical procedures were carried out with approved protocols at the University of Michigan. The electrode array was inserted by hand using #5 fine PTFE-coated forceps into the target cortical area. Typically, the electrode was inserted such that the top site was just below the cortical surface. The electrode assembly was wrapped with GelFoam (Pfizer, Inc., New York) and then cemented with dental acrylic. The skin around the acrylic was tightened with sutures and anti-bacterial cream was applied. Animals were given 3-4 days to recover post-surgery before experiments were resumed.

4.2.2 Recording Procedure

After recovery from the surgery, animals continued to perform the same behavioral task while neural activity was recorded from both hemispheres. Spike times, spike waveforms, LFPs, and external events were recorded simultaneously using a Plexon Multichannel Acquisition Processor (MAP, Plexon Inc., Dallas, TX). The signal from each electrode was passed through a high-input impedance headstage with unity gain and then filtered by a preamplifier to extract the spike and the LFP components as follows: LFPs were obtained by applying a filter with a passband of 3Hz to 500Hz and were sampled and digitized at 1kHz. Spike signals were filtered with a passband of 154-8800Hz, further amplified, and sampled at 40kHz. All behavioral events and relevant timestamps were also recorded and were sampled at 25μs resolution. Data were recorded in a single session each day, for a period of 4-6 weeks.

4.2.3 Determination of electrode array locations

Details of the procedure used to determine the electrode array locations are described in Section 3.2.4. Briefly, at the end of the study, we performed microlesions on sites at top, middle, and bottom of each electrode array by passing 35μA DC for 2s using a potentiostat (AUTOLAB, Eco Chemie, Netherlands). We waited 2-3 hours for microlesion 'scars' to form, and then perfused the animal intracardially with 4% paraformaldehyde and explanted the brain tissue. After fixing the tissue, we took 50μm coronal sections and then performed a standard cresyl-violet Nissl stain on the tissue slices. The slices were then analyzed under a microscope and images were taken to reconstruct the position of the electrode array based on the shank track, lesion marks, and the known geometry of the probe. We identified the location of the boundary between the upper and lower layers (the start of Layer 5) by analyzing the gray-level index values of the images using ImageJ (RSB, NIH: http://rsb.info.nih.gov/ij/). Fig. 3.3 shows array tracks, electrode lesion marks, and the boundary between the upper and lower cortical layers in brain slices from all animals (D2-D4) used in the subsequent analysis.

4.2.4 Post-processing, filtering and spectral analysis

Trials were aligned by behavioral epochs and then cut into analysis windows. Only correct behavioral trials were used for the subsequent analysis. The noisy trials were detected and removed as follows: after computing the mean and standard deviation (std), all trials that exceeded the $\pm 1.96 \times std$ bounds were removed. Data was then analyzed in four frequency bands – low (3-15Hz), beta-gamma (15-40Hz), high gamma (40-70Hz) and high frequency (70-150Hz).

The LFP signals were separated in the various frequency bands using digital filters

implemented in MATLAB. Low (3-15Hz) and high (70-150Hz) bands were filtered by using a 9^{th} order Butterworth digital filter, (`butter` function in MATLAB). Butterworth filters were also used for bandpass filtering. The order for the bandpass filters was obtained using `butterord` function after specifying the pass-bands of interest for high-gamma (passband between 40-70Hz, stopband frequencies, Fs(low) = 20Hz and Fs(high) = 100Hz) and beta-gamma (passband between 15-40Hz, stopband frequencies, Fs(low) = 5Hz and Fs(high) = 70Hz). The filters were designed with 5dB ripple in the passband and 40dB attenuation in the stopband. To eliminate any phase shifts, signals were filtered forward and backward using the `filtfilt` function. Leading and trailing portions were discarded to account for filtering artifacts.

4.3 Results

This study is based on 585 LFP recordings from 39 sessions from three rats (D2-D4) with implants in both hemispheres. Figure 3.3 shows the location of the six electrode arrays; LFPs were recorded across all the different layers.

4.3.1 Evoked potentials across the different layers in different frequency ranges

All correct trials were separated by movement towards the left or right, and were aligned to the movement onset. Evoked potentials were analyzed in the four different frequency bands as described earlier. Figure 4.1 shows evoked potentials for movement to each direction from a typical unit in the different frequency bands of interest across all layers. In the low frequency band, activity was mostly flat till after the onset of movement. Differences in direction of movement towards the left and right were most apparent in the low frequency band. Figure 4.2 shows spectrograms from all six implants across the different layers around movement onset for frequency range from 5 to 100Hz. Most of the directional activity was observed

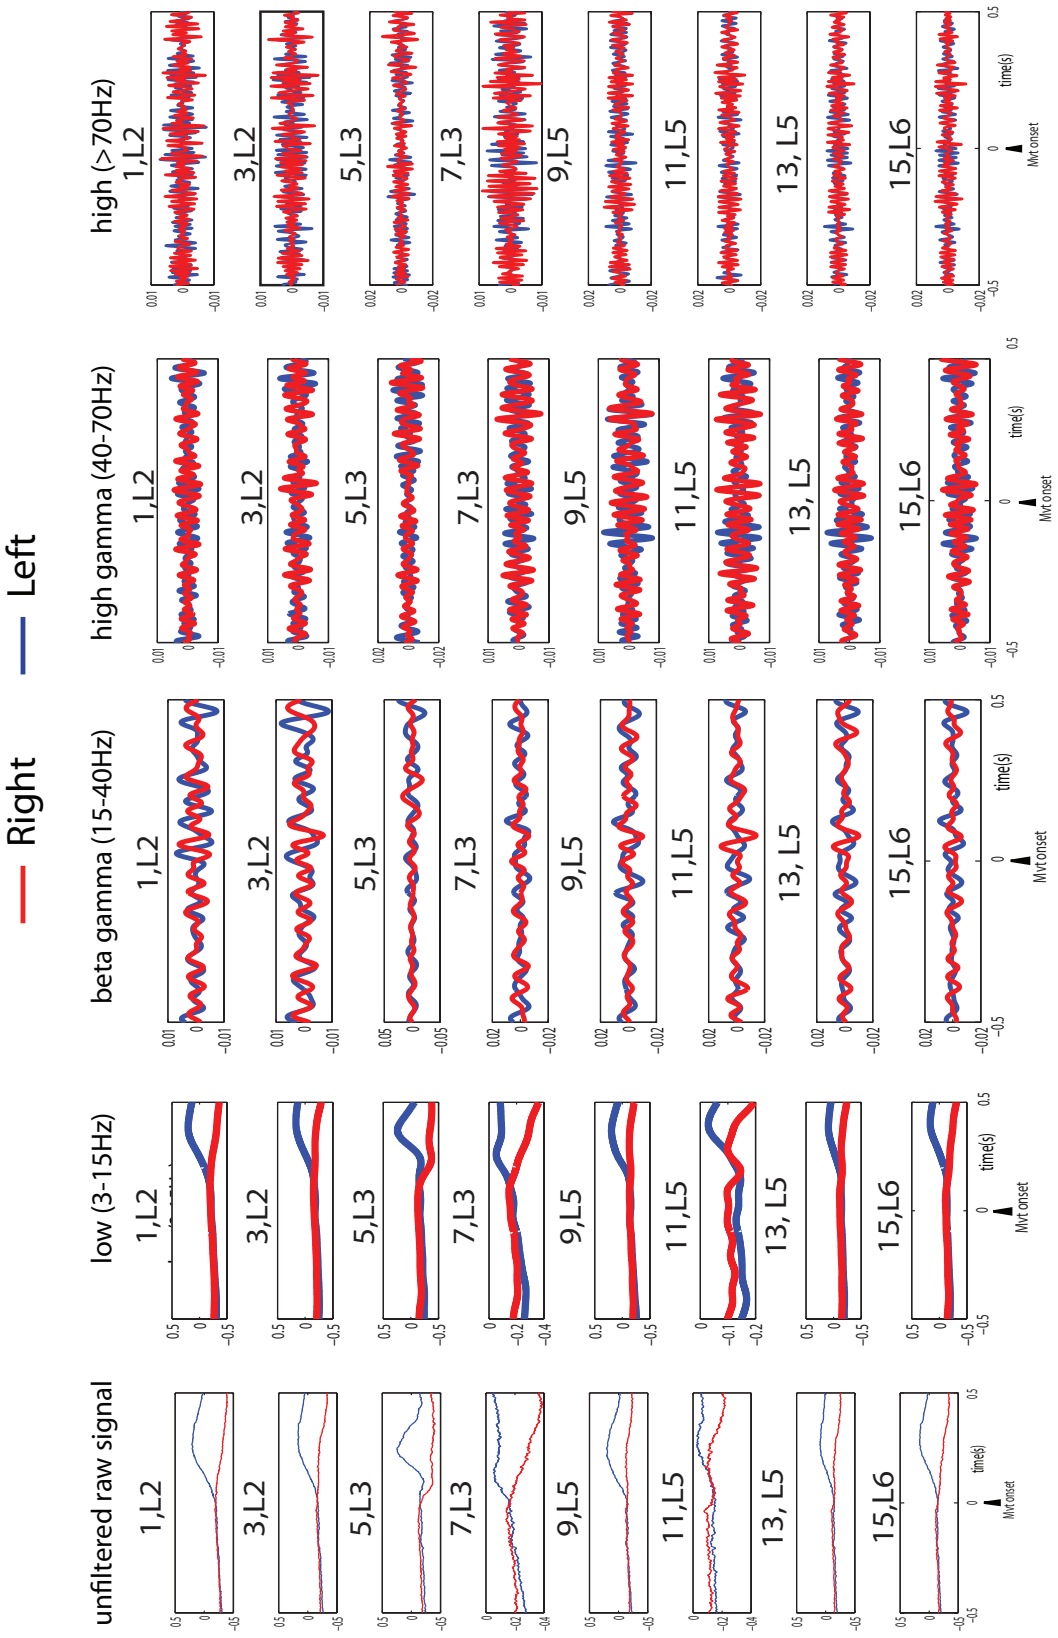

Figure 4.1: Average evoked potentials (EP) around movement onset for the unfiltered data and in the different frequency bands shown separately for eight sites spanning across all layers for one sample implant. Data is for all correct trials and shows EPs for rightward and leftward trials. The EP for the low frequency band (3-15Hz) shows a distinct difference between movements in the two directions. Short epochs of oscillations are observed in the high-gamma and high frequency bands.

in the low frequency band. It has been previously observed that LFPs were slower than units in indicating the direction of movement [Asher et al., 2007]. We confirm this late activation and further add that this mainly occurs in the low frequency band. No discernible phase-locked oscillations or decrease in activity were seen in the beta-gamma band. In the high-gamma band, phase-locked oscillatory activity was observed $\approx$ 120 to 180ms before the onset of movement. This oscillatory activity was predominant in the lower layers (Layers 5,6). Phase-locked oscillatory activity was also seen in the high frequency band across all layers.

Figure 4.3 shows evoked potentials in the different frequency ranges with activity from all layers shown on one plot for one typical unit in one particular direction. In the low frequency range, LFPs from sites located in the Layer 5 and Layer 2 were almost identical. Evoked potential activity in Layer 3 was often found to be quite distinct from activity in the other layers. Activity in the beta-gamma range was quite similar across all layers for a considerable portion of the analysis epoch. As shown in the middle panel in Figure 4.3, activity in the higher gamma range was synchronized across all layers only for a small interval ($\approx$ 200ms) around movement onset. Thereafter, only activity in the lower Layers 5,6 was synchronized. Activity in the high frequency range was synchronized for shorter periods, being more similar in phase across layers $\approx$ 200ms before movement onset. Other shorter periods of activity similar in phase lasting less than 100ms were also noticed.

4.3.2 Encoding of direction in the different frequency ranges

Direction analysis was performed to determine variation in the LFP with respect to leftward or rightward movement in the different layers in each frequency band. For each LFP trial, the mean RMS amplitude was computed for the 1s time window centered around movement onset. The evoked potentials in each direction were then

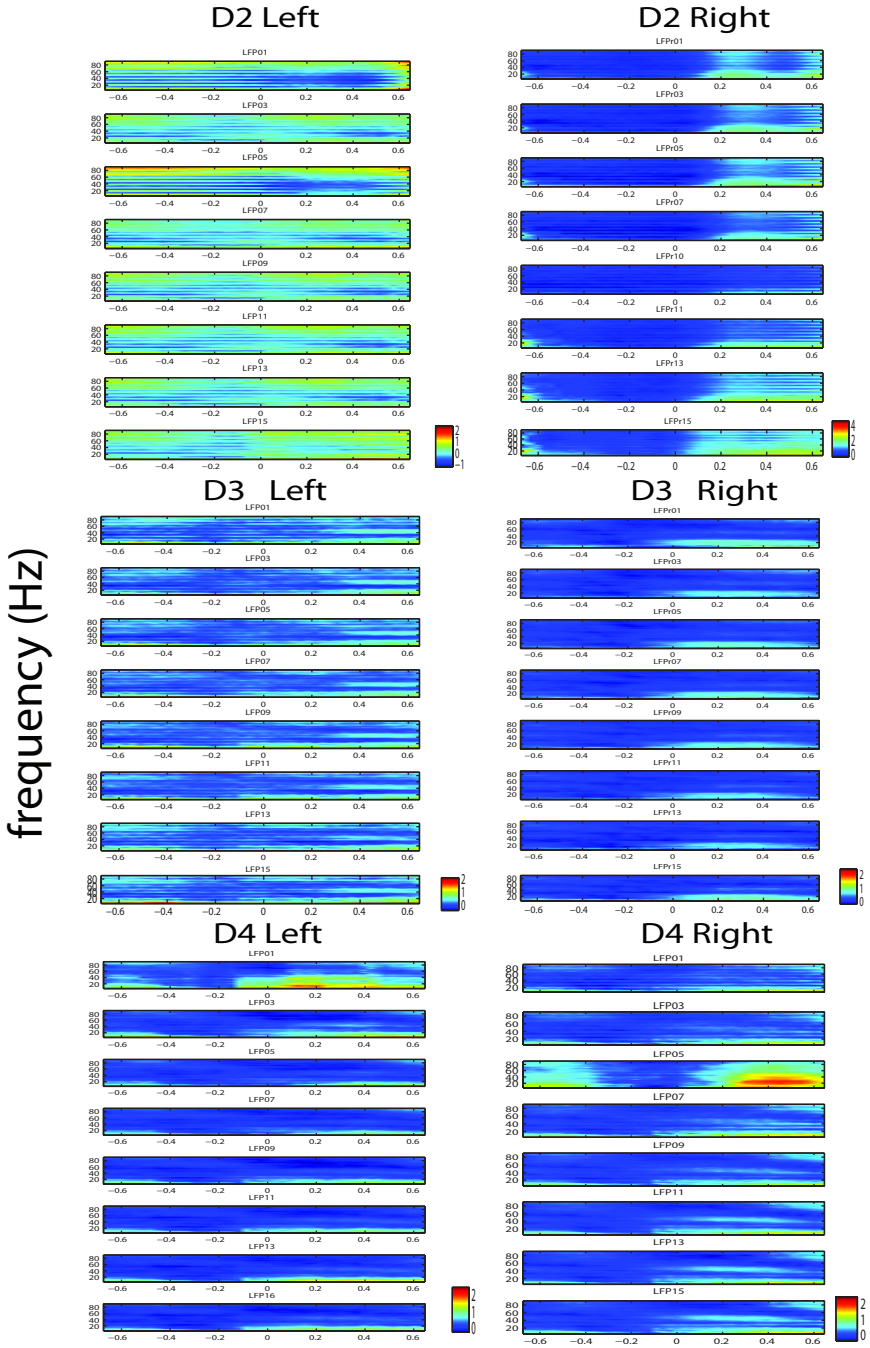

Figure 4.2: Averaged spectrogram across layers for one sample session for movement towards the contralateral direction. Spectrogram values shown on log scale, for frequencies between 5-100Hz, t=0 denotes movement onset, LFP channels arranged in depthwise with channel numbers corresponding to the layers indicated in 1.1. Across all six animals, activity is chiefly concentrated in the low frequency band. While, increased activity was observed in particular frequency bands in some cases; overall, there was no consistent difference that observed in frequencies above 15Hz.

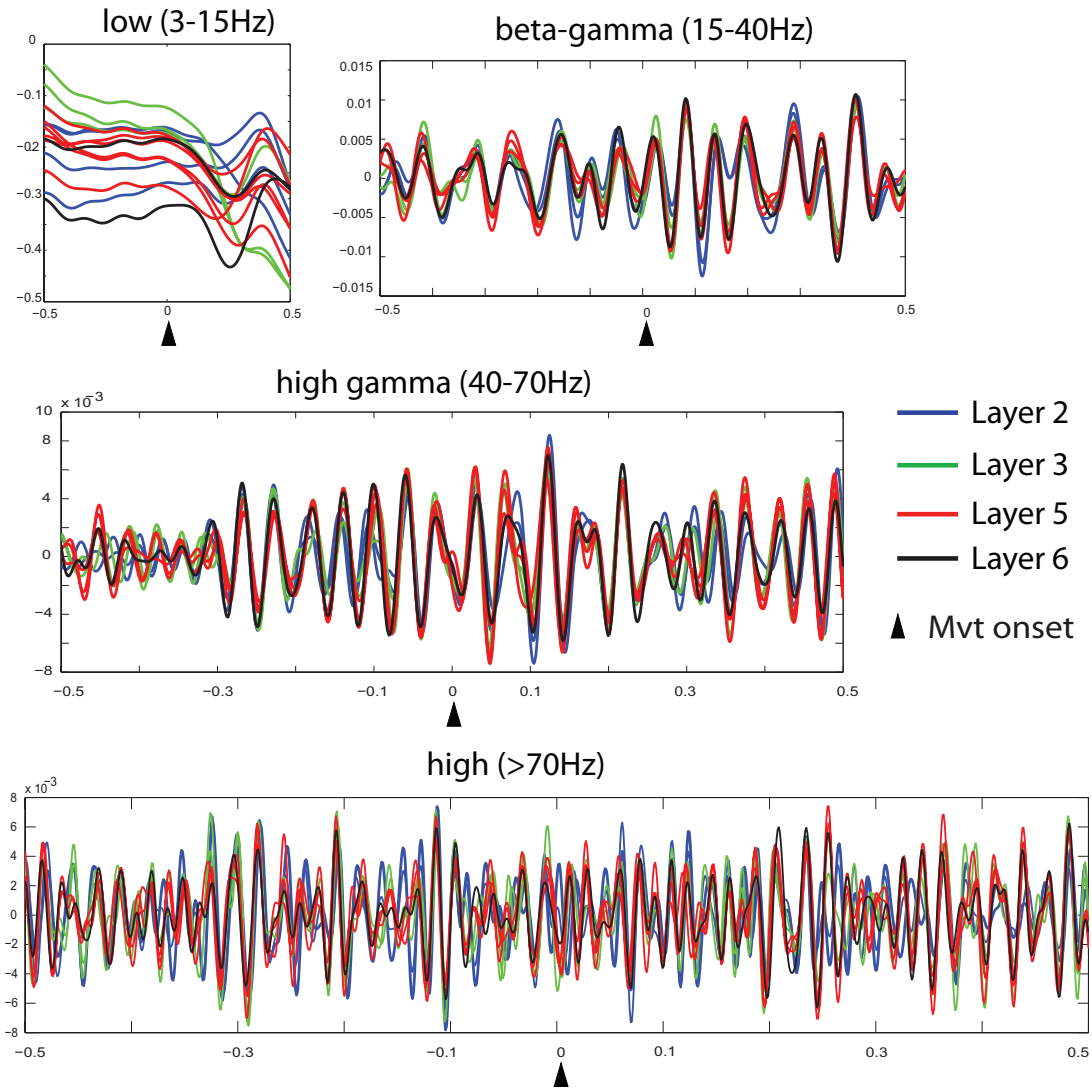

Figure 4.3: Average evoked potentials from different layers in the different frequency bands: Typical plot of average evoked potentials from one animal showing activity in the different bands for one particular direction of movement for LFPs from all the different layers (color coded as shown). Only alternate sites are plotted for the purpose of visualization. In the low frequency band, the activity differences occur late and similarity in phase between LFPs is low. The highest similarity in phase was detected in the beta-gamma (15-40Hz) band. The high-gamma (40-70Hz) band showed similarity in phase between the LFPs in the different layers only around the onset of movement. Short epochs of similarity in phase were observed in high frequency (>70Hz) band. LFPs from the lower layers shown in L5 (red) and L6 (black) tended to be more coherent with each other.

compared using a Wilcoxon ranksum test to determine if activity was significantly different. Figure 4.4 shows the proportion of LFPs that significantly encoded either left or rightward direction in the various frequency bands for each of the six implants considered separately. We compared performance between each of the bands using a 2x2 contingency table. A Bonferroni correction was applied to correct for the multiple comparisons. Overall, the low frequency band showed significantly different encoding performance compared to other bands. Performance in the other frequency bands was statistically indistinct from each other. As the Figure 4.4 shows there is variation in performance across the different implants. In at least two implants D2L and D4L, the beta-gamma band showed performance that was statistically indistinct from direction encoding in the low frequency band.

When we compared direction encoding between the upper and lower layers, we found a statistically significant difference only in the high-gamma frequency range. As Figure 4.5 demonstrates, in all other frequency ranges the fraction of LFPs encoding direction were approximately the same. Thus, excepting the high-gamma band, we fail to reject the null hypothesis that LFP activity is the same with respect to direction encoding between the upper and lower layers of the motor cortex.

4.4 Discussion

4.4.1 Direction encoding in the different frequency bands

In our study, we observed that across all implants, LFPs in the low frequency band (3-15Hz) contained the most direction information. The direction information was not exclusively contained in the low-frequency band, as directional encoding was observed in some fraction of the LFPs in all frequency bands. In a few cases, the fraction of LFPs in the beta-gamma (15-40Hz) and the high-gamma (40-70Hz) bands encoded directional information that was indistinct from low frequency band. Our

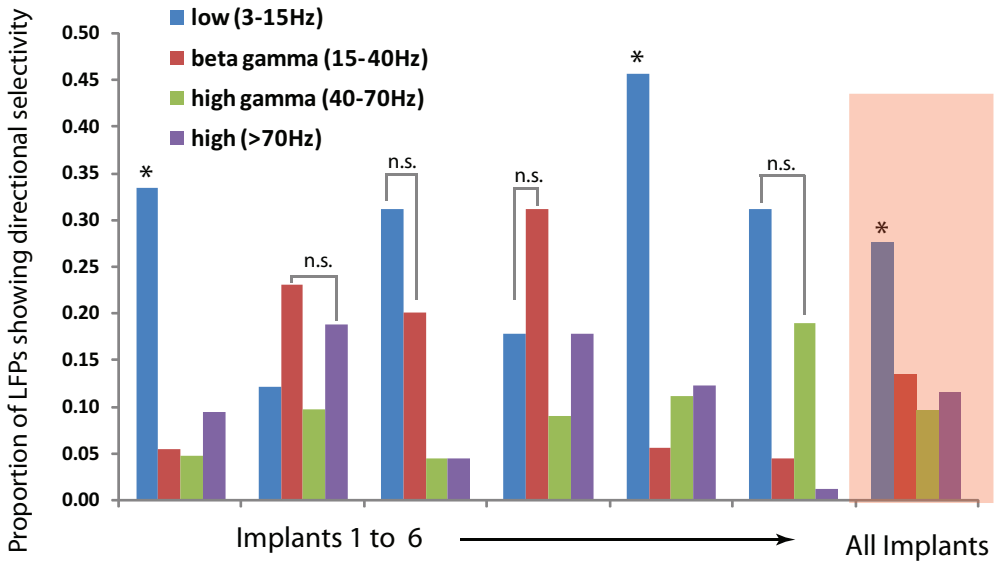

Figure 4.4: Histograms show the fraction of LFPs in each of the frequency bands that showed significant difference in activity in either direction. Plots are shown for each of the six implants and for all implants considered together. Pairs of frequency bands were tested for significance. Asterisk indicates that particular group was significantly different from the other groups. A Bonferroni correction was applied to correct for multiple comparisons across groups.

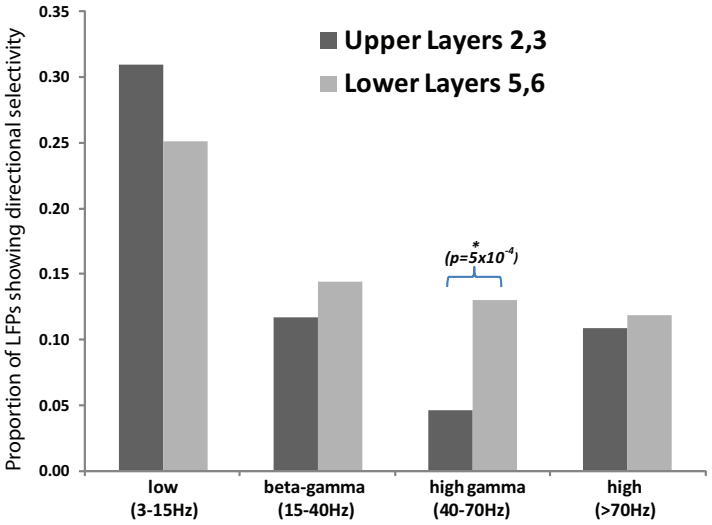

Figure 4.5: Proportion of LFPs showing significant direction encoding in the upper and lower layers with respect to RMS amplitude for either leftward or rightward movement in each of the frequency bands: low (3-15Hz), beta-gamma(15-40Hz), high-gamma(40-70Hz) and high (70-150Hz).

results are similar to another study that found the best directional tuning in the low (<10Hz) frequency band, and weak decoding in the fast (25-40Hz) and intermediate (10-25Hz) frequency bands [O'Leary and Hatsopoulos, 2006]. We also confirm results from [Rickert et al., 2005] that found that movement evoked potentials (MEPs) in the <4Hz and 6-13Hz band had the best decoding ability. Contrary to their finding, we found that activity in the 16-42Hz also provides directional information; however, performance in this band was not uniform across animals. Overall, our results suggest that the fraction of LFPs in the low frequency bands contain as much (or better) direction information relative to other frequency bands. Directional information can be extracted from the higher frequency bands and can be combined with the low frequency band resulting in better overall decoding of direction as shown by [Rickert et al., 2005]. LFPs have shown selective activation in the gamma band in response to attention the visual [Womelsdorf et al., 2006] and parietal areas of the brain [Pesaran and Movshon, 2008]. Additional experiments would need to be performed in a multidirectional, 3-D task to verify if the different LFP frequency bands encode independent movement parameter information such as velocity, acceleration, joint dynamics, etc.

4.4.2 Direction encoding across the different layers

LFPs are occurring due to synaptic or neural activity within 250μm of the recording electrode [Katzner et al., 2009], and not on the order of millimeters as suggested earlier [Mitzdorf, 1985]. Hence, our observation of similarity in phase of LFPs across the different layers of the cortex is not simply due to a common origin for these signals. It has also been confirmed that the brain acts as a pure resistor and thus transmits both high and low frequencies equally well [Logothetis et al., 2007]. We observed short-term similarity in phase across layers in the high-gamma and

high frequency bands. Researchers shown increased coherence in specific frequencies between brain areas during specific behavioral tasks [Bekisz and Wrbel, 1999], [Liang et al., 2002]. These findings have led to the hypothesis that coherence in the local field potentials between brain regions is a method of cortical communication [Fries, 2005].

Our analysis of activity within layers confirm previous observations that have shown that LFPs recorded at short inter-electrode distances are highly redundant [Legatt et al., 1980], and within-area LFPs showed high correlations in selectivity [Spinks et al., 2008]. Indeed, activity in the lower layers (Layers 5,6) was more correlated to each other than to LFP activity in the upper layers. It has been hypothesized that the upper layers (Layers 2,3) perform feature selection after integrating sub-cortical and inter-areal input, while the lower layers (Layers 5,6) are involved in actual output [Douglas and Martin, 2004]. A recent study verified aspects of this theory by showing that when upper layer neurons in the M1 are stimulated, network wide events are evoked implying that upper layer neurons drive output neurons in lower layers [Weiler et al., 2008]. In our study, except in the high-gamma band, there was no distinguishable difference between the upper and lower layers. Thus, control information from LFPs can be obtained from either upper or layers of the cortex for comparable performance. Since LFPs are locally more correlated, from a neuroprosthetic perspective, more spatially distant sampling is suggested, i.e. LFPs should be sampled from electrode sites that are more widely spaced, or collected from different shanks for more independent information.

4.4.3 Comparison between unit activity and LFP activity

In comparison to our earlier study, the performance from LFPs was comparable to that obtained from single units for the given analysis window cen-

tered around movement onset. It was shown earlier that by combining the activity of LFPs and single units (SUAs) or LFPs and multi-units(MUAs), a superior prediction accuracy was obtained as compared to that provided by just using these signals alone [Mehring et al., 2003]. Signals conveying different cognitive functions are useful for creating multiple channels for communication [Andersen et al., 2004]. The extent of specificity of LFPs remains to be investigated. Studies have shown that while information can be extracted from LFPs, single units have shown pre-movement activity that encodes the upcoming movement [Achtman et al., 2007],[Hatsopoulos et al., 1998] and evoked potentials have shown to lag behind single units [Asher et al., 2007]. This implies that neuroprostheses based on using LFPs alone may not be appropriate for rapidly changing movements. Also, LFPs do not have uniform distribution of directional preferences, and thus do not convey as much information as single-units which have a more distributed directional preference [O'Leary and Hatsopoulos, 2006], [Scherberger et al., 2005]. Dynamic coherence that was phase-locked with respect to movement onset in the different frequency bands suggests that oscillations in the LFPs may play a role in communication or coordinating the activity of neurons in the different layers [Fries, 2005]. Whether this oscillatory structure can be exploited for a neuroprosthetic control signals is still to be determined. Given the advantages of long-term stability and reliability LFPs continue to remain a viable alternative cortical signals, but oscillatory behavior and the role of the different frequency bands across the lamina needs to be investigated further.

CHAPTER V

Conclusion

If you can dream – and not make dreams your master,

If you can think – and not make thoughts your aim;
If you can meet with Triumph and Disaster
And treat those two impostors just the same;
...
If you can fill the unforgiving minute
With sixty seconds' worth of distance run,
Yours is the Earth and everything that's in it,
And–which is more–you'll be a Man, my son!
- Rudyard Kipling, *If*

Rapid progress has been made in the last decade in the arena of neuroprostheses. Devices and systems have increased in sophistication and recent work is increasingly meeting the realistic demands for such devices. The first human implant vindicated the efforts of the past decade using animal models and demonstrated cursor control. The device was controlled exclusively using brain signals and operated simulated e-mail and devices such as a television. Furthermore, control was shown to open and close a prosthetic hand, and perform basic tasks [Hochberg et al., 2006]. Recent work showed a monkey feeding itself using a mechanical arm – cortical signals controlled a gripper on the end of the arm and demonstrated physical interaction between the monkey, the robotic arm and objects in the workspace [Velliste et al., 2008]. The next level of sophistication has been reached in demonstrating effective closed-loop

control of neuroprostheses. Recently, closed-loop direct control of paralyzed muscles by cortical neurons was demonstrated to restore volitional control to paralyzed limbs via electrical stimulation by creating artificial connections between the cortex and muscles in monkeys [Moritz et al., 2008].

The first human implant was a technological and scientific achievement [Hochberg et al., 2006], but long-term efficacy was limited as the recordings did not last beyond ten months, causing the patient to abandon the use of the device. A fully-functional, reliable, long-term human neuroprosthesis requires progress in three critical areas: a) methods for interpreting and extracting control information from a variety of signals, b) algorithms for reliable and effective single-trial decoding, and c) advances in the neuron-electrode interface. This dissertation focused on only one of the above aspects, i.e. improving the effectiveness of unit and LFP activity using chronic microelectrodes. The problem of obtaining appropriate control signals along with the most effective modality is far from solved, but parallel improvements in neural interface technology are also absolutely critical.

5.1 Future directions of dissertation studies

5.1.1 Improvements in modeling and metrics

In Chapter II, we demonstrate a novel method for tracking and assessing neural signatures across sessions. The model is first trained on data from one session; the regression parameters obtained are then used to predict similar units from spike clusters obtained from two different sessions of recording. The logistic regression model as presented assumes that the regression parameters in the model, i.e. βs, are constant. Due to intrinsic and extrinsic factors, neural signatures of the same unit are rarely stationary across sessions. There are drifts in metric space due to changes at the neural interface due to tissue micromotion [Subbaroyan et al., 2005], and glial

encapsulation. There are also changes in the intrinsic neural signature, such as firing patterns, due to plasticity and learning. Our method can be further improved by incorporating these non-stationary effects by accounting for these drifts across time in the regression model. One way to achieve this would be by using a mixed effects model. A mixed effects model includes both fixed and random factors and allows the parameters to vary within a certain bound.

The standard linear regression model is given by:

$$y_i = \beta_1 x_{1i} + \beta_2 x_{2i} + \ldots + \beta_n x_{ni} + \epsilon_i, \text{ where } \epsilon_i \sim \mathcal{N}(0, \sigma^2) \tag{5.1}$$

which has has only one random effect, the error term ϵ_i. The parameters of the model are the regression coefficients, $\beta_1, \beta_2, \ldots, \beta_n$.

Mixed-effect models include additional random-effect terms, which are appropriate for representing time-dependent data. A factor is random if the effects associated with the levels of the factor can be viewed as being like a random sample from a population of effects. For random effects, we can make statements about variation in the population of random effects from which the effects at hand are considered to be like a random sample. Furthermore, we can generalize our conclusions about fixed factors to the populations associated with random factors.

To incorporate mixed effects we can modify Equation 5.1 as,

$$y_{ij} = \beta_1 x_{1ij} + \ldots + \beta_p x_{pij} + b_{i1} z_{1ij} + \ldots + b_{iq} x_{qij} + \epsilon_{ij} \tag{5.2}$$

where,

$$\begin{aligned} b_{ik} &\sim \mathcal{N}(0, \psi^2) \\ Cov(b_k, b_{k'}) &= \psi_{kk'} \\ \epsilon_i &\sim \mathcal{N}(0, \sigma^2 \lambda_{ijj}) \\ Cov(\epsilon_{ij}, \epsilon_{ij'}) &= \sigma^2 \lambda_{ijj'} \end{aligned}$$

In our study, we studied only a limited set of nine metrics. Our method only provides a framework for probabilistic assessment, the actual efficacy of the method depends on appropriately chosen metrics. Depending on the nature of the experiment, animal preparation (awake or anaesthetized), location and types of cells recorded, metrics should be modified or new metrics adopted. Continuously recorded data would help determine the effects of the neuronal drift more accurately and also help verify our method. It has been shown that drifts on a short timescale may require adaptive signal processing methodologies [Linderman et al., 2008]. There are many different spike sorting methods that can be tweaked to perform tracking under the framework as proposed.

5.1.2 Improved layer and cell labeling techniques

In Chapter III, we showed that units in the lower layers of the cortex are more likely to encode movement and direction information than units in the upper layers. Better localization of sites in the individual layers will help determine more subtle differences in activity [Naselaris et al., 2006]. Using functional and electrophysiological criteria, two kinds of putative pyramidal cells were found in the motor cortex: PP1 cells which were broadly tuned and located in all layers of the cortex, and PP2 cells that were narrowly tuned [Merchant et al., 2008]. The directional tuning of PP1 cells was strongly affected by dynamic sculpting via putative interneurons. While determining electrode site locations is important, advances in techniques to determine the location and type of cell recorded in an awake, behaving preparation via extracellular recordings will help validate proposed cortical microcircuits and the functional role of the different cells across the various layers [Du et al., 2008].

The task used in our study was a two-direction task. In a future experiment, incorporating more directions, measuring additional movement parameters such as

velocity, acceleration, grip, etc. will enable us to more realistically model the requirements for a neuroprosthesis and parse finer details of movement.

5.1.3 Local field potentials as alternative sources for cortical signals

In Chapter IV, we showed that LFPs contain directional information in different bands and our results indicate that the low frequency bands contain the most information. Our results suggest that the different frequency bands of LFPs contain a rich set of information that needs to be further investigated in a more complicated 3-D task. Except in the high-gamma band, our experiment was unable to determine a significant difference between activity in the upper versus lower layers in terms of direction encoding. The similarity of phase observed in our data may be related to the dynamic coherence that has been shown to be phase-locked with respect to behavioral events in the different frequency bands. Whether this oscillatory and phase structure can be exploited for a neuroprosthetic control signals is still to be determined.

The correlations among the LFPs within layers and the local origin of these signals suggest that densely packed electrodes may not be able to obtain adequate decoding performance from LFPs for a neuroprosthesis [O'Leary and Hatsopoulos, 2006]. LFPs and units were shown to encode different information even on same electrode [Asher et al., 2007]. Joint analysis of units with LFPs needs further development. The question of how LFPs fit within the framework of cortical microcircuits is still unresolved.

REFERENCES

[Achtman et al., 2007] Achtman, N., Afshar, A., Santhanam, G., Yu, B. M., Ryu, S. I., and Shenoy, K. V. (2007). Free-paced high-performance brain-computer interfaces. *J Neural Eng*, 4(3):336–347.

[Albright et al., 1984] Albright, T. D., Desimone, R., and Gross, C. G. (1984). Columnar organization of directionally selective cells in visual area mt of the macaque. *J Neurophysiol*, 51(1):16–31.

[Amirikian and Georgopoulos, 2003] Amirikian, B. and Georgopoulos, A. P. (2003). Modular organization of directionally tuned cells in the motor cortex: is there a short-range order? *Proc Natl Acad Sci U S A*, 100(21):12474–12479.

[Andersen et al., 2004] Andersen, R. A., Musallam, S., and Pesaran, B. (2004). Selecting the signals for a brain-machine interface. *Curr Opin Neurobiol*, 14(6):720–726.

[Aoki et al., 1999] Aoki, F., Fetz, E. E., Shupe, L., Lettich, E., and Ojemann, G. A. (1999). Increased gamma-range activity in human sensorimotor cortex during performance of visuomotor tasks. *Clin Neurophysiol*, 110(3):524–537.

[Asher et al., 2007] Asher, I., Stark, E., Abeles, M., and Prut, Y. (2007). Comparison of direction and object selectivity of local field potentials and single units in macaque posterior parietal cortex during prehension. *J Neurophysiol*, 97(5):3684–3695.

[Baker et al., 1999] Baker, S. N., Kilner, J. M., Pinches, E. M., and Lemon, R. N. (1999). The role of synchrony and oscillations in the motor output. *Exp Brain Res*, 128(1-2):109–117.

[Baker et al., 1997] Baker, S. N., Olivier, E., and Lemon, R. N. (1997). Coherent oscillations in monkey motor cortex and hand muscle EMG show task-dependent modulation. *J Physiol*, 501 (Pt 1):225–241.

[Barth et al., 2004] Barth, P., Hirase, H., Monconduit, L., Zugaro, M., Harris, K. D., and Buzski, G. (2004). Characterization of neocortical principal cells and interneurons by network interactions and extracellular features. *J Neurophysiol*, 92(1):600–608.

[Bauer et al., 1983] Bauer, R., Dow, B. M., Snyder, A. Z., and Vautin, R. (1983). Orientation shift between upper and lower layers in monkey visual cortex. *Exp Brain Res*, 50(1):133–145.

[Bekisz and Wrbel, 1999] Bekisz, M. and Wrbel, A. (1999). Coupling of beta and gamma activity in corticothalamic system of cats attending to visual stimuli. *Neuroreport*, 10(17):3589–3594.

[Benabid et al., 2001] Benabid, A. L., Koudsie, A., Benazzouz, A., Piallat, B., Krack, P., Limousin-Dowsey, P., Lebas, J. F., and Pollak, P. (2001). Deep brain stimulation for Parkinson's disease. *Adv Neurol*, 86:405–412.

[Bennett and Lemon, 1996] Bennett, K. M. and Lemon, R. N. (1996). Corticomotoneuronal contribution to the fractionation of muscle activity during precision grip in the monkey. *J Neurophysiol*, 75(5):1826–1842.

[Berke et al., 2004] Berke, J. D., Okatan, M., Skurski, J., and Eichenbaum, H. B. (2004). Oscillatory entrainment of striatal neurons in freely moving rats. *Neuron*, 43(6):883–896.

[Brecht et al., 2004] Brecht, M., Krauss, A., Muhammad, S., Sinai-Esfahani, L., Bellanca, S., and Margrie, T. W. (2004). Organization of rat vibrissa motor cortex and adjacent areas according to cytoarchitectonics, microstimulation, and intracellular stimulation of identified cells. *J Comp Neurol*, 479(4):360–373.

[Britten et al., 1996] Britten, K. H., Newsome, W. T., Shadlen, M. N., Celebrini, S., and Movshon, J. A. (1996). A relationship between behavioral choice and the visual responses of neurons in macaque mt. *Vis Neurosci*, 13(1):87–100.

[Buzsaki, 2004] Buzsaki, G. (2004). Large-scale recording of neuronal ensembles. *Nat Neurosci*, 7(5):446–451.

[Chestek et al., 2007] Chestek, C. A., Batista, A. P., Santhanam, G., Yu, B. M., Afshar, A., Cunningham, J. P., Gilja, V., Ryu, S. I., Churchland, M. M., and Shenoy, K. V. (2007). Single-neuron stability during repeated reaching in macaque premotor cortex. *J Neurosci*, 27(40):10742–10750.

[Cohen and Nicolelis, 2004] Cohen, D. and Nicolelis, M. A. L. (2004). Reduction of single-neuron firing uncertainty by cortical ensembles during motor skill learning. *J Neurosci*, 24(14):3574–3582.

[Cooper, 1981] Cooper, I. S. (1981). Twenty-five years of experience with physiological neurosurgery. *Neurosurgery*, 9(2):190–200.

[Csicsvari et al., 2003] Csicsvari, J., Henze, D. A., Jamieson, B., Harris, K. D., Sirota, A., Bartho, P., Wise, K. D., and Buzsaki, G. (2003). Massively Parallel Recording of Unit and Local Field Potentials With Silicon-Based Electrodes. *J Neurophysiol*, 90(2):1314–1323.

[DeFelipe et al., 2002] DeFelipe, J., Alonso-Nanclares, L., and Arellano, J. I. (2002). Microstructure of the neocortex: comparative aspects. *J Neurocytol*, 31(3-5):299–316.

[Delescluse and Pouzat, 2006] Delescluse, M. and Pouzat, C. (2006). Efficient spike-sorting of multi-state neurons using inter-spike intervals information. *J Neurosci Methods*, 150(1):16–29.

[Delgado, 1967] Delgado, J. M. (1967). Aggression and defense under cerebral radio control. *UCLA Forum Med Sci*, 7:171–193.

[Delgado, 1969] Delgado, J. M. (1969). Conditioned suppression via subcortical radio stimulation in the chimpanzee. *Tech Doc Rep ARL TDR*, 1:1–16.

[Donchin et al., 2001] Donchin, O., Gribova, A., Steinberg, O., Bergman, H., de Oliveira, S. C., and Vaadia, E. (2001). Local field potentials related to bimanual movements in the primary and supplementary motor cortices. *Exp Brain Res*, 140(1):46–55.

[Donoghue, 2002] Donoghue, J. P. (2002). Connecting cortex to machines: recent advances in brain interfaces. *Nat Neurosci*, 5 Suppl:1085–1088.

[Donoghue et al., 1998] Donoghue, J. P., Sanes, J. N., Hatsopoulos, N. G., and Gal, G. (1998). Neural discharge and local field potential oscillations in primate motor cortex during voluntary movements. *J Neurophysiol*, 79(1):159–173.

[Donoghue and Wise, 1982] Donoghue, J. P. and Wise, S. P. (1982). The motor cortex of the rat: cytoarchitecture and microstimulation mapping. *J Comp Neurol*, 212(1):76–88.

[Douglas and Martin, 2004] Douglas, R. J. and Martin, K. A. C. (2004). Neuronal circuits of the neocortex. *Annu Rev Neurosci*, 27:419–451.

[Du et al., 2008] Du, J., Riedel-Kruse, I. H., Nawroth, J. C., Roukes, M. L., Laurent, G. J., and Masmanidis, S. C. (2008). High-resolution three-dimensional extracellular recording of neuronal activity with microfabricated electrode arrays. *J Neurophysiol*, page http://dx.doi.org/10.1152/jn.90992.2008.

[Emondi et al., 2004] Emondi, A. A., Rebrik, S. P., Kurgansky, A. V., and Miller, K. D. (2004). Tracking neurons recorded from tetrodes across time. *J Neurosci Methods*, 135(1-2):95–105.

[Faraway, 2006] Faraway, J. J. (2006). *Extending the Linear Model with R*, chapter Binomial Data, pages 25–54. Chapman & Hall.

[Fee et al., 1996] Fee, M. S., Mitra, P. P., and Kleinfeld, D. (1996). Automatic sorting of multiple unit neuronal signals in the presence of anisotropic and non-gaussian variability. *J Neurosci Methods*, 69(2):175–188.

[Firth, 1993] Firth, D. (1993). Bias reduction of maximum likelihood estimates. *Biometrika*, 80(1):27–38.

[Fries, 2005] Fries, P. (2005). A mechanism for cognitive dynamics: neuronal communication through neuronal coherence. *Trends in Cognitive Sciences*, 9(10):474 – 480.

[Fries et al., 2001] Fries, P., Reynolds, J. H., Rorie, A. E., and Desimone, R. (2001). Modulation of oscillatory neuronal synchronization by selective visual attention. *Science*, 291(5508):1560–1563.

[Georgopoulos et al., 1989] Georgopoulos, A. P., Crutcher, M. D., and Schwartz, A. B. (1989). Cognitive spatial-motor processes. 3. Motor cortical prediction of movement direction during an instructed delay period. *Exp Brain Res*, 75(1):183–194.

[Georgopoulos et al., 2007] Georgopoulos, A. P., Merchant, H., Naselaris, T., and Amirikian, B. (2007). Mapping of the preferred direction in the motor cortex. *Proc Natl Acad Sci U S A*, 104(26):11068–11072.

[Gilja et al., 2006] Gilja, V., Linderman, M. D., Santhanam, G., Afshar, A., Ryu, S., Meng, T. H., and Shenoy, K. V. (2006). Multiday electrophysiological recordings from freely behaving primates. *Conf Proc IEEE Eng Med Biol Soc*, 1:5643–5646.

[Harris et al., 2000] Harris, K. D., Henze, D. A., Csicsvari, J., Hirase, H., and Buzski, G. (2000). Accuracy of tetrode spike separation as determined by simultaneous intracellular and extracellular measurements. *J Neurophysiol*, 84(1):401–414.

[Hastie et al., 2001] Hastie, T., Tibshirani, R., and Friedman, J. H. (2001). *The Elements of Statistical Learning*. Springer.

[Hatsopoulos et al., 2004] Hatsopoulos, N., Joshi, J., and O'Leary, J. G. (2004). Decoding continuous and discrete motor behaviors using motor and premotor cortical ensembles. *J Neurophysiol*, 92(2):1165–1174.

[Hatsopoulos et al., 1998] Hatsopoulos, N. G., Ojakangas, C. L., Paninski, L., and Donoghue, J. P. (1998). Information about movement direction obtained from synchronous activity of motor cortical neurons. *Proc Natl Acad Sci U S A*, 95(26):15706–15711.

[Hetke JF, 2002] Hetke JF, A. D. (2002). *In Handbook of Neuroprosthetic Methods*, chapter Silicon microelectrodes for extracellular recording, page 7: 16391. CRC, Boca Raton, FL.

[Hochberg et al., 2006] Hochberg, L. R., Serruya, M. D., Friehs, G. M., Mukand, J. A., Saleh, M., Caplan, A. H., Branner, A., Chen, D., Penn, R. D., and Donoghue, J. P. (2006). Neuronal ensemble control of prosthetic devices by a human with tetraplegia. *Nature*, 442(7099):164–171.

[Humphrey et al., 1970] Humphrey, D. R., Schmidt, E. M., and Thompson, W. D. (1970). Predicting Measures of Motor Performance from Multiple Cortical Spike Trains. *Science*, 170(3959):758–762.

[Hutsler et al., 2005] Hutsler, J. J., Lee, D.-G., and Porter, K. K. (2005). Comparative analysis of cortical layering and supragranular layer enlargement in rodent carnivore and primate species. *Brain Res*, 1052(1):71–81.

[Katzner et al., 2009] Katzner, S., Nauhaus, I., Benucci, A., Bonin, V., Ringach, D. L., and Carandini, M. (2009). Local origin of field potentials in visual cortex. *Neuron*, 61(1):35–41.

[Kennedy et al., 2000] Kennedy, P. R., Bakay, R. A., Moore, M. M., Adams, K., and Goldwaithe, J. (2000). Direct control of a computer from the human central nervous system. *IEEE Trans Rehabil Eng*, 8(2):198–202.

[Kipke et al., 2008] Kipke, D. R., Shain, W., Buzski, G., Fetz, E., Henderson, J. M., Hetke, J. F., and Schalk, G. (2008). Advanced neurotechnologies for chronic neural interfaces: new horizons and clinical opportunities. *J Neurosci*, 28(46):11830–11838.

[Krimer et al., 2005] Krimer, L. S., Zaitsev, A. V., Czanner, G., Krner, S., Gonzlez-Burgos, G., Povysheva, N. V., Iyengar, S., Barrionuevo, G., and Lewis, D. A. (2005). Cluster analysis-based physiological classification and morphological properties of inhibitory neurons in layers 2-3 of monkey dorsolateral prefrontal cortex. *J Neurophysiol*, 94(5):3009–3022.

[Legatt et al., 1980] Legatt, A. D., Arezzo, J., and Vaughan, H. G. (1980). Averaged multiple unit activity as an estimate of phasic changes in local neuronal activity: effects of volume-conducted potentials. *J Neurosci Methods*, 2(2):203–217.

[Leuthardt et al., 2004] Leuthardt, E. C., Schalk, G., Wolpaw, J. R., Ojemann, J. G., and Moran, D. W. (2004). A brain–computer interface using electrocorticographic signals in humans. *Journal of Neural Engineering*, 1(2):63–71.

[Lewicki, 1998] Lewicki, M. S. (1998). A review of methods for spike sorting: the detection and classification of neural action potentials. *Network*, 9(4):R53–R78.

[Liang et al., 2002] Liang, H., Bressler, S. L., Ding, M., Truccolo, W. A., and Nakamura, R. (2002). Synchronized activity in prefrontal cortex during anticipation of visuomotor processing. *Neuroreport*, 13(16):2011–2015.

[Linderman et al., 2008] Linderman, M. D., Santhanam, G., Kemere, C. T., Gilja, V., O'Driscoll, S., Yu, B. M., Afshar, A., Ryu, S. I., Shenoy, K. V., and Meng, T. H. (2008). Signal processing challenges for neural prostheses. *IEEE Signal Processing Magazine*, 25(1):18–28.

[Liu et al., 2006] Liu, X., McCreery, D. B., Bullara, L. A., and Agnew, W. F. (2006). Evaluation of the stability of intracortical microelectrode arrays. *IEEE Trans Neural Syst Rehabil Eng*, 14(1):91–100.

[Liu et al., 1999] Liu, X., McCreery, D. B., Carter, R. R., Bullara, L. A., Yuen, T. G., and Agnew, W. F. (1999). Stability of the interface between neural tissue and chronically implanted intracortical microelectrodes. *IEEE Trans Rehabil Eng*, 7(3):315–326.

[Loeb, 1990] Loeb, G. E. (1990). Cochlear prosthetics. *Annual Review of Neuroscience*, 13(1):357–371.

[Logothetis et al., 2007] Logothetis, N. K., Kayser, C., and Oeltermann, A. (2007). In vivo measurement of cortical impedance spectrum in monkeys: implications for signal propagation. *Neuron*, 55(5):809–823.

[Marzullo et al., 2006] Marzullo, T. C., Miller, C. R., and Kipke, D. R. (2006). Suitability of the cingulate cortex for neural control. *IEEE Trans Neural Syst Rehabil Eng*, 14(4):401–409.

[Maynard et al., 1999] Maynard, E. M., Hatsopoulos, N. G., Ojakangas, C. L., Acuna, B. D., Sanes, J. N., Normann, R. A., and Donoghue, J. P. (1999). Neuronal Interactions Improve Cortical Population Coding of Movement Direction. *J. Neurosci.*, 19(18):8083–8093.

[McCreery et al., 2006] McCreery, D., Lossinsky, A., Pikov, V., and Liu, X. (2006). Microelectrode array for chronic deep-brain microstimulation and recording. *IEEE Trans Biomed Eng*, 53(4):726–737.

[McLachlan, 1997] McLachlan, R. S. (1997). Vagus nerve stimulation for intractable epilepsy: a review. *J Clin Neurophysiol*, 14(5):358–368.

[McNaughton et al., 1983] McNaughton, B. L., O'Keefe, J., and Barnes, C. A. (1983). The stereotrode: a new technique for simultaneous isolation of several single units in the central nervous system from multiple unit records. *J Neurosci Methods*, 8(4):391–397.

[Mehring et al., 2003] Mehring, C., Rickert, J., Vaadia, E., de Oliveira, S. C., Aertsen, A., and Rotter, S. (2003). Inference of hand movements from local field potentials in monkey motor cortex. *Nat Neurosci*, 6(12):1253–1254.

[Merchant et al., 2008] Merchant, H., Naselaris, T., and Georgopoulos, A. P. (2008). Dynamic sculpting of directional tuning in the primate motor cortex during three-dimensional reaching. *J Neurosci*, 28(37):9164–9172.

[Miller et al., 2007] Miller, K. J., Leuthardt, E. C., Schalk, G., Rao, R. P. N., Anderson, N. R., Moran, D. W., Miller, J. W., and Ojemann, J. G. (2007). Spectral changes in cortical surface potentials during motor movement. *J Neurosci*, 27(9):2424–2432.

[Mitzdorf, 1985] Mitzdorf, U. (1985). Current source-density method and application in cat cerebral cortex: investigation of evoked potentials and eeg phenomena. *Physiol Rev*, 65(1):37–100.

[Moritz et al., 2008] Moritz, C. T., Perlmutter, S. I., and Fetz, E. E. (2008). Direct control of paralysed muscles by cortical neurons. *Nature*, 456(7222):639–642.

[Mountcastle, 1997] Mountcastle, V. B. (1997). The columnar organization of the neocortex. *Brain*, 120 (Pt 4):701–722.

[Mountcastle, 2003] Mountcastle, V. B. (2003). Introduction. computation in cortical columns. *Cereb Cortex*, 13(1):2–4.

[Mountcastle et al., 1957] Mountcastle, V. B., Davies, P. W., and Berman, A. L. (1957). Response properties of neurons of cat's somatic sensory cortex to peripheral stimuli. *J Neurophysiol*, 20(4):374–407.

[Muir and Steeves, 1997] Muir, G. D. and Steeves, J. D. (1997). Sensorimotor stimulation to improve locomotor recovery after spinal cord injury. *Trends Neurosci*, 20(2):72–77.

[Murthy and Fetz, 1992] Murthy, V. N. and Fetz, E. E. (1992). Coherent 25- to 35-hz oscillations in the sensorimotor cortex of awake behaving monkeys. *Proc Natl Acad Sci U S A*, 89(12):5670–5674.

[Murthy and Fetz, 1996] Murthy, V. N. and Fetz, E. E. (1996). Oscillatory activity in sensorimotor cortex of awake monkeys: synchronization of local field potentials and relation to behavior. *J Neurophysiol*, 76(6):3949–3967.

[Naselaris et al., 2005] Naselaris, T., Merchant, H., Amirikian, B., and Georgopoulos, A. P. (2005). Spatial reconstruction of trajectories of an array of recording microelectrodes. *J Neurophysiol*, 93(4):2318–2330.

[Naselaris et al., 2006] Naselaris, T., Merchant, H., Amirikian, B., and Georgopoulos, A. P. (2006). Large-scale organization of preferred directions in the motor cortex. i. motor cortical hyperacuity for forward reaching. *J Neurophysiol*, 96(6):3231–3236.

[Nelson et al., 2008] Nelson, M. J., Pouget, P., Nilsen, E. A., Patten, C. D., and Schall, J. D. (2008). Review of signal distortion through metal microelectrode recording circuits and filters. *J Neurosci Methods*, 169(1):141–157.

[Neves and Ruther, 2007] Neves, H. P. and Ruther, P. (2007). The neuroprobes project. *Conf Proc IEEE Eng Med Biol Soc*, 2007:6443–6445.

[Nicolelis, 2001] Nicolelis, M. A. (2001). Actions from thoughts. *Nature*, 409(6818):403–407.

[Nicolelis et al., 2003] Nicolelis, M. A. L., Dimitrov, D., Carmena, J. M., Crist, R., Lehew, G., Kralik, J. D., and Wise, S. P. (2003). Chronic, multisite, multielectrode recordings in macaque monkeys. *Proc Natl Acad Sci U S A*, 100(19):11041–11046.

[Nordhausen et al., 1996] Nordhausen, C. T., Maynard, E. M., and Normann, R. A. (1996). Single unit recording capabilities of a 100 microelectrode array. *Brain Res*, 726(1-2):129–140.

[O'Leary and Hatsopoulos, 2006] O'Leary, J. G. and Hatsopoulos, N. G. (2006). Early visuomotor representations revealed from evoked local field potentials in motor and premotor cortical areas. *J Neurophysiol*, 96(3):1492–1506.

[Otto et al., 2005] Otto, K. J., Rousche, P. J., and Kipke, D. R. (2005). Cortical microstimulation in auditory cortex of rat elicits best-frequency dependent behaviors. *J Neural Eng*, 2(2):42–51.

[Parikh et al., 2009] Parikh, H., Marzullo, T., and Kipke, D. (2009). Lower layers in the motor cortex are efficient targets for penetrating microelectrodes in cortical prostheses. *Journal of Neural Engineering*, 6:026004.

[Patil and Turner, 2008] Patil, P. G. and Turner, D. A. (2008). The development of brain-machine interface neuroprosthetic devices. *Neurotherapeutics*, 5(1):137–146.

[Pesaran and Movshon, 2008] Pesaran, B. and Movshon, J. A. (2008). What to do, or how to do it? *Neuron*, 58(3):301–303.

[Pesaran et al., 2008] Pesaran, B., Nelson, M. J., and Andersen, R. A. (2008). Free choice activates a decision circuit between frontal and parietal cortex. *Nature*, 453(7193):406–409.

[Pesaran et al., 2002] Pesaran, B., Pezaris, J. S., Sahani, M., Mitra, P. P., and Andersen, R. A. (2002). Temporal structure in neuronal activity during working memory in macaque parietal cortex. *Nat Neurosci*, 5(8):805–811.

[Pfurtscheller et al., 2003] Pfurtscheller, G., Graimann, B., Huggins, J. E., Levine, S. P., and Schuh, L. A. (2003). Spatiotemporal patterns of beta desynchronization and gamma synchronization in corticographic data during self-paced movement. *Clin Neurophysiol*, 114(7):1226–1236.

[Rennaker et al., 2005] Rennaker, R. L., Street, S., Ruyle, A. M., and Sloan, A. M. (2005). A comparison of chronic multi-channel cortical implantation techniques: manual versus mechanical insertion. *J Neurosci Methods*, 142(2):169–176.

[Rickert et al., 2005] Rickert, J., de Oliveira, S. C., Vaadia, E., Aertsen, A., Rotter, S., and Mehring, C. (2005). Encoding of Movement Direction in Different Frequency Ranges of Motor Cortical Local Field Potentials. *J. Neurosci.*, 25(39):8815–8824.

[Rockland and Ichinohe, 2004] Rockland, K. S. and Ichinohe, N. (2004). Some thoughts on cortical minicolumns. *Exp Brain Res*, 158(3):265–277.

[Rokni et al., 2003] Rokni, U., Steinberg, O., Vaadia, E., and Sompolinsky, H. (2003). Cortical representation of bimanual movements. *J Neurosci*, 23(37):11577–11586.

[Salinas and Romo, 1998] Salinas, E. and Romo, R. (1998). Conversion of sensory signals into motor commands in primary motor cortex. *J Neurosci*, 18(1):499–511.

[Samejima et al., 2005] Samejima, K., Ueda, Y., Doya, K., and Kimura, M. (2005). Representation of action-specific reward values in the striatum. *Science*, 310(5752):1337–1340.

[Santhanam et al., 2007] Santhanam, G., Linderman, M. D., Gilja, V., Afshar, A., Ryu, S. I., Meng, T. H., and Shenoy, K. V. (2007). Hermesb: a continuous neural recording system for freely behaving primates. *IEEE Trans Biomed Eng*, 54(11):2037–2050.

[Santhanam et al., 2006] Santhanam, G., Ryu, S. I., Yu, B. M., Afshar, A., and Shenoy, K. V. (2006). A high-performance brain-computer interface. *Nature*, 442(7099):195–198.

[Schalk et al., 2007] Schalk, G., Kubnek, J., Miller, K. J., Anderson, N. R., Leuthardt, E. C., Ojemann, J. G., Limbrick, D., Moran, D., Gerhardt, L. A., and Wolpaw, J. R. (2007). Decoding two-dimensional movement trajectories using electrocorticographic signals in humans. *J Neural Eng*, 4(3):264–275.

[Schalk et al., 2008] Schalk, G., Miller, K. J., Anderson, N. R., Wilson, J. A., Smyth, M. D., Ojemann, J. G., Moran, D. W., Wolpaw, J. R., and Leuthardt, E. C. (2008). Two-dimensional movement control using electrocorticographic signals in humans. *J Neural Eng*, 5(1):75–84.

[Scherberger et al., 2005] Scherberger, H., Jarvis, M. R., and RA, A. (2005). Cortical local field potential encodes movement intentions in the posterior parietal cortex. *Neuron*, 46:347–54.

[Schwartz, 2004] Schwartz, A. B. (2004). Cortical neural prosthetics. *Annu Rev Neurosci*, 27:487–507.

[Serruya et al., 2002] Serruya, M. D., Hatsopoulos, N. G., Paninski, L., Fellows, M. R., and Donoghue, J. P. (2002). Instant neural control of a movement signal. *Nature*, 416(6877):141–142.

[Seymour and Kipke, 2007] Seymour, J. P. and Kipke, D. R. (2007). Neural probe design for reduced tissue encapsulation in cns. *Biomaterials*, 28(25):3594–3607.

[Shoham et al., 2001] Shoham, S., Halgren, E., Maynard, E. M., and Normann, R. A. (2001). Motor-cortical activity in tetraplegics. *Nature*, 413(6858):793.

[Snider and Bonds, 1998] Snider, R. K. and Bonds, A. B. (1998). Classification of non-stationary neural signals. *J Neurosci Methods*, 84(1-2):155–166.

[Spinks et al., 2008] Spinks, R. L., Kraskov, A., Brochier, T., Umilta, M. A., and Lemon, R. N. (2008). Selectivity for grasp in local field potential and single neuron activity recorded simultaneously from m1 and f5 in the awake macaque monkey. *J Neurosci*, 28(43):10961–10971.

[Subbaroyan et al., 2005] Subbaroyan, J., Martin, D. C., and Kipke, D. R. (2005). A finite-element model of the mechanical effects of implantable microelectrodes in the cerebral cortex. *J Neural Eng*, 2(4):103–113.

[Suner et al., 2005] Suner, S., Fellows, M. R., Vargas-Irwin, C., Nakata, G. K., and Donoghue, J. P. (2005). Reliability of signals from a chronically implanted, silicon-based electrode array in non-human primate primary motor cortex. *IEEE Trans Neural Syst Rehabil Eng*, 13(4):524–541.

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

[Taylor et al., 2003] Taylor, D. M., Tillery, S. I. H., and Schwartz, A. B. (2003). Information conveyed through brain-control: cursor versus robot. *IEEE Trans Neural Syst Rehabil Eng*, 11(2):195–199.

[Tolias et al., 2007] Tolias, A. S., Ecker, A. S., Siapas, A. G., Hoenselaar, A., Keliris, G. A., and Logothetis, N. K. (2007). Recording chronically from the same neurons in awake, behaving primates. *J Neurophysiol*, 98(6):3780–3790.

[Velliste et al., 2008] Velliste, M., Perel, S., Spalding, M. C., Whitford, A. S., and Schwartz, A. B. (2008). Cortical control of a prosthetic arm for self-feeding. *Nature*, 453(7198):1098–1101.

[Vetter et al., 2004] Vetter, R. J., Williams, J. C., Hetke, J. F., Nunamaker, E. A., and Kipke, D. R. (2004). Chronic neural recording using silicon-substrate microelectrode arrays implanted in cerebral cortex. *IEEE Trans Biomed Eng*, 51(6):896–904.

[Weiler et al., 2008] Weiler, N., Wood, L., Yu, J., Solla, S. A., and Shepherd, G. M. G. (2008). Top-down laminar organization of the excitatory network in motor cortex. *Nat Neurosci*, 11(3):360–366.

[Williams et al., 1999] 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 Res Brain Res Protoc*, 4(3):303–313.

[Wise et al., 2004] Wise, K. D., Anderson, D. J., Hetke, J. F., Kipke, D. R., and Najafi, K. (2004). Wireless implantable microsystems: high-density electronic interfaces to the nervous system. *Proceedings of the IEEE*, 92(1):76–97.

[Wolpaw, 2004] Wolpaw, J. R. (2004). Brain-computer interfaces (BCIs) for communication and control: a mini-review. *Suppl Clin Neurophysiol*, 57:607–613.

[Wolpaw and McFarland, 2004] Wolpaw, J. R. and McFarland, D. J. (2004). Control of a two-dimensional movement signal by a noninvasive brain-computer interface in humans. *Proc Natl Acad Sci U S A*, 101(51):17849–17854.

[Womelsdorf et al., 2006] Womelsdorf, T., Fries, P., Mitra, P. P., and Desimone, R. (2006). Gamma-band synchronization in visual cortex predicts speed of change detection. *Nature*, 439(7077):733–736.