Responses of Multiple Simultaneously Recorded Macaque Area LIP Neurons in a Memory Saccade Task

Thesis by
John Stylianos Pezaris

In Partial Fulfillment of the Requirements
for the Degree of
Doctor of Philosophy

1891

California Institute of Technology
Pasadena, California
2000
(Submitted May 18, 2000)

Again, to my parents, Dr. Pezaris and Dr. Pezaris.

Acknowledgments

Above all, I must thank my collaborator, Maneesh Sahani, without whom this thesis would be meager and pitiful. Together, our complementary talents formed a substantial alliance, and our work has placed us at the forefront of exciting research. Despite what Maneesh might have to eloquently say on the subject, I suspect that mine was the better end of the deal.

I would like to thank my advisor, Richard Andersen, for providing opportunities remarkable in their generosity, and for supporting our foray into the technological and scientific unknown that has proved rich and rewarding. My committee, Gilles Laurent, Christof Koch, Erin Schumann, and Pietro Perona, have in turn, listened to problems, asked the right questions, and guided me towards solutions, for which I am grateful. I would like to thank Gilles in particular for supporting the collaboration with Mike Wehr which appears in Chapter 7.

For an unbounded supply of moral and emotional support, I would like to especially acknowledge Holli Weld. She has been remarkable and tireless in trying to provide everything.

My friends deserve much credit for keeping me on the academic straight and narrow by providing a personal outlet that strayed well beyond. Thanks to Desmond Patterson, Donal Gallagher, Stephanie McCarthy, Bob Nichols, and Adam Kent, who were there with a wealth of companionship. My physically more distant friends, John Wolfe, Bill, Lori and Emma Goldenthal, Sharon Thomas, Juli Zimmerman, Nicho Hatsopoulos, Tatiana Flessas, Sean Hatsopoulos, Dave Goddeau, Mike Blair, and Dottie Dartland, all did their part, knowingly or not.

Our support staff, Cierina Reyes, Betty Gillikin, and Viktor Shcherbatyuk have kept things working smoothly despite the many hiccoughs that happen in an experimental laboratory.

Current and former members of the Andersen laboratory have provided intellectual support and good friendship. Thanks, in no particular order, to Flip Sabes and Alexander Grunewald who never tired of my naive statistical questions; Daniella Meeker who helped with everything; David Dubowitz who had the crazy idea to put a monkey in a magnet; Yale Cohen who made surgery fun; David Bradley who taught me how to cut; Ken Grieve who did the histology presented in Chapter 3, and was an all-round jolly fellow; Krishna Shenoy who was a superlative intellectual sounding board; Bjorn Christianson who helped me figure out why I, repeatedly, was being stupid; Larry Snyder who inspired by example; Martyn Bracewell who gave valued advice and knowledge; and Murray Jarvis who lended a hand whenever asked.

The enthusiasm with which people offered their assistance in preparing drafts of this dissertation was touching. I am indebted to my parents, Daniella Meeker, Holli Weld, Murray Jarvis, and David Dubowitz. The document is far better for their input and assistance.

> *Folks, thanks so much for encouraging my behavior.*
> — NICK CARTER, Signoff banter on WFNX (1993)

Responses of Multiple Simultaneously Recorded Macaque Area LIP Neurons in a Memory Saccade Task

John Stylianos Pezaris

Submitted to the Computation and Neural Systems Program
on November 1, 1999, in partial fulfillment of the requirements for
the degree of Doctor of Philosophy in Computation and Neural Systems

Abstract

Cells in the lateral intraparietal cortex (LIP) of rhesus macaques respond vigorously and in spatially-tuned fashion to briefly memorized visual stimuli. Responses to stimulus presentation, memory maintenance, and task completion are seen, in varying combination from neuron to neuron. To help elucidate this functional segmentation a new system for simultaneous recording from multiple neighboring neurons was developed. The two parts of this dissertation discuss the technical achievements and scientific discoveries, respectively.

Technology. Simultanous recordings from multiple neighboring neurons were made with four-wire bundle electrodes, or tetrodes, which were adapted to the awake behaving primate preparation. Signals from these electrodes were partitionable into a background process with a $1/f$-like spectrum and foreground spiking activity spanning 300–6000 Hz. Continuous voltage recordings were sorted into spike trains using a state-of-the-art clustering algorithm, producing a mean of 3 cells per site. The algorithm classified 96% of spikes correctly when tetrode recordings were confirmed with simultaneous intracellular signals. Recording locations were verified with a new technique that creates electrolytic lesions visible in magnetic resonance imaging, eliminating the need for histological processing. In anticipation of future multi-tetrode work, the *chronic chamber microdrive*, a device for long-term tetrode delivery, was developed.

Science. Simultaneously recorded neighboring LIP neurons were found to have similar preferred targets in the memory saccade paradigm, but dissimilar peristimulus time histograms, PSTH). A majority of neighboring cell pairs had a difference in preferred directions of under 45° while the trial time of maximal response showed a broader distribution, suggesting homogeneity of tuning with heterogeneity of function. A continuum of response characteristics was present, rather than a set of specific response types; however, a mapping experiment suggests this may be because a given cell's PSTH changes shape as well as amplitude through the response field. Spike train autocovariance was tuned over target and changed through trial epoch, suggesting different mechanisms during memory versus background periods. Mean frequency-domain spike-to-spike coherence was concentrated below 50 Hz with a significant maximum of 0.08; mean time-domain coherence had a narrow peak in the range ±10 ms with a significant maximum of 0.03. Time-domain coherence was found to be untuned for short lags (10 ms), but significantly tuned at larger lags (50 ms).

Thesis Advisor: Richard A. Andersen
Title: James G. Boswell Professor of Neuroscience

Table Of Contents

Table Of Tables

Table Of Figures

I don't dispute the doubts you've outlined
But it's my right to waste your time
— SOUL COUGHING (Fully Retractable, *El Oso*, 1998)

Chapter One: Introduction

1 Introduction

Exploitation of locality is a guiding principle of design, from temporal or spatial locality in computer memory systems, to multi-processor communications, to built architecture, urban design, and biological systems. In particular, what has come to be known in Electrical Engineering as *the wiring constraint*, is found in areas as disparate as roadway design, space module engineering, and the analysis of neural systems. While it is this last area that forms the basis for this dissertation, the wiring constraint is perhaps best expressed as the observation that in modern integrated circuits, less than one-tenth of the chip area is devoted to active silicon, the remainder being required for wires to transfer signals from one transistor to another; because of this overwhelming dominance of wiring needs, designers take great care to place related components in close proximity. Placing components adjacent to related elements optimizes many things, including total system resources, speed, efficiency of power distribution and waste heat extraction. This is true at all levels of system design, from transistor, to subcircuit, module, chip and board.

Given this driving force observed in man-made systems, one is led to ask whether complex biological systems follow the same principle. The answer is yes, again, at many levels; for example, it guides the formation of tissue with like specificity into organs, and parts of these organs according to subfunction. In the mammalian brain, we find a segmentation of tissue comprised of neurons and the axonal processes communicating between them into grey matter around the periphery and white matter in the center, respectively. We find gray matter organized into distinct layers, each with particular functions, such as input, output, or local connectivity. On top of this laminar structure, we find neurons with certain functions to be grouped together into areas, and areas arranged in a hierarchical fashion, such as shown in van Essen's diagram of the primate visual system [18]. The visual field is optically mapped on to the retina, where each light-sensitive rod and cone cell responds to light in one small pixel-like part of the visual space. This partitioning is carried through to the higher cortical areas where we continue to find spatial locality. For the areas thus-far explored, cells sensitive to one part of space at one processing stage project directly to cells at the next level sensitive to the same part of space, as shown for example by Reid and Alonso for thalamic to cortical visual projections [61].

We thus are led to the central question behind this dissertation: do neurons which are physically adjacent communicate? Are they part of a subcircuit, such as a column or microcolumn, interacting mostly with neighbors, or do they only communicate with distant cells? Within the guise of artificial neural networks, this is akin to asking if networks are purely feedforward or contain recurrent links. While our vast experience as an industry in designing electronics circuitry would suggest the latter will certainly be discovered to be true, it remains an open question as to the full character and extent of local versus non-local connectivity in brain circuitry.

To understand a circuit, we need to monitor all components within the circuit. Given the necessarily limited ability to make such measurements, it is useful to monitor components in local groups, rather than in disparate areas (although that is useful as well, but for other purposes). To this end we investigated the adaptation of a relatively new method for simultaneously recording from multiple neighboring neurons, the bundled microwire electrode. This proved successful, and we were able to routinely record from 3–5 physically adjacent cells with four-wire electrodes called *tetrodes*.

2 Location: LIP

We chose to make these recordings in the dorsal visual stream. While it might have been advisable to approach a more fully understood area such as primary visual cortex when adapting tetrodes to awake behaving primate use and in the exploration of interactions between neighboring neurons, the research interests of our laboratory tend to higher order areas of visual processing, and in particular the parietal cortex. Given the varying efficacies of training and recording data from areas subsuming different sensory and behavioral modalities, selecting a part of the visual pathway as a sphere for study, and therefore using a visual task to drive the experiments, is unassailable. Visual tasks are easy to train and produce vigorous responses from the visual cortices. Recording from the parietal cortex allows us to investigate interesting behaviors that contain a cognitive element, and thus require alert animals. In particular, much work has been done in our laboratory studying the lateral bank of the intraparietal sulcus, or area LIP, and while the area's physical location within the brain makes delivery of electrodes and some interpretation of the results more difficult, the ability to explore a relatively newly identified part of cortex is exciting.

Lesions of the posterior parietal cortex in humans and monkeys impair spatial perception, and deleteriously affect saccade generation [6, 26, 36, 38]. The area is bilaterally located a little posterior to the stereotaxic zero plane which passes through the ear canals, approximately halfway up the skull from the ears to the top of the head. This location, between visual and motor areas, is reflected in the projections to and from the area, placing it at the cusp of visual sensory-motor integration.

From studies by Mountcastle [43] and Andersen [4], we know that posterior parietal neuronal activity includes visual and saccade-related components, that is, when a target is visually presented to a subject, neurons respond both to the target presentation as well as during the subsequent gaze shift when the subject visually acquires the target. These first studies in parietal cortex, part of Ungerleider and Desimone's *where* pathway [5] subserving spatial perception, struck controversy as saccade-related activity was, in two camps, alternately ascribed to the motor command of saccade generation [43], or a purely sensory response to the visual stimulus [65]. Later work [4] teased apart the visual and motor components by using an experimental paradigm which temporally dissociated the stimulus presentation and saccade execution. In this *memory saccade* task, originally introduced by Hikosaka and Wurtz [27], the animal is required to temporarily memorize the location of a flashed target, postponing a saccade to the target for approximately one second's delay until the fixation light is extinguished. Using this task it was discovered [21] that in addition to sensory and motor responses, now separated in time, some neurons also responded during the delay between target flash and saccade execution, despite the absence of a visual target. These neurons, it was argued, displayed the trace of a working or short-term spatial memory system.

Anatomical evidence was then used to identify a subregion of the parietal cortex on the lateral bank of the sulcus called the lateral intraparietal area, or LIP [37, 2, 9]. Responses in this area to the memory saccade have been mapped using the memory saccade in various guises, including single and double saccades [3, 40, 7, 8], and have been shown to be spatially tuned. A given cell will have a response profile that varies as a stimulus is placed in different parts of visual field, with a concentrated area of activity in a well-defined part of visual space called the *receptive field*. The same cell will possibly also have a *motor field* that describes the responses associated with gaze shifts to varying points in the space. These two fields generally share a common maximum point and have similar shapes, all relative to the fixation point, although, to complicate matters, the magnitude of the response is often scaled by a factor which depends upon the current gaze position in a planar way from one extreme of the visual field to another [3]. Additional work [7] mapped out the *memory field*, or the response versus target position during the memory period, and found that at least in coarse form, the memory fields of active neurons had strong similarity to the visual and motor fields. Thus LIP neurons were found to have three phases of response to the memory saccade task, visual, memory, and saccadic, which, when present, would have overlapping tuning; but not all cells were found to have all phases.

3 Evidence for Homogeneity and Heterogeneity

Implications from previous results suggest that nearby neurons do communicate strongly with each other. From work such as by Georgopoulos in motor cortex [33], Donoghue in premotor cortex [25], Vaadia in auditory cortex [1], Gerstein in inferotempral cortex [22], and many others, we know that communication between adjacent or nearby cells can be found ubiquitously in the brain, and that the interactions vary with stimulus condition.

Do these neurons then share similar characteristics, or are they different? Previous results point in both directions. Starting with Hubel and Wiesel's original work [29], there was evidence for a smooth transitioning of characteristics from cell to cell, suggesting that neighboring neurons have similar responses. Later, work with optical imaging in primary visual cortex [10] suggested that neurons with similar characteristics tend to cluster, and even that like cells project to like [61]. In parietal cortex, where we have been working, receptive fields across the area were initially found to have smooth transitions bounded by abrupt fragmentary borders [9], in line with the other evidence suggesting a homogeneity of responses at the local scale.

On the other hand, there is evidence such as from hippocampus where neighboring neurons have little overlap in characteristics with their neighbors [73], or from primary visual cortex where receptive fields have higher scatter from cell to cell than previously thought [39], or somatosensory cortex with similar findings [66]. All these are suggestive of a heterogeneity of responses at the local scale.

In turn, we sought to examine this question in parietal cortex at the local level, between physically adjacent neurons, determining whether such cells were similar or dissimilar in response characteristics. Other reports on LIP show that distinct phases of activity through a memory saccade task exist, not all of which are present in every cell, suggestive of a heterogeneity [9, 7, 8]. For example, while Mazzoni's primary finding [40] is that in LIP most neurons encode the next intended saccade, even if there is more than one planned, there is a small, but significant, minority that respond to a given saccade plan no matter if there are others planned or not. These two populations might be expected to co-mingle to generate a heterogeneity of function. Our findings, presented in this thesis, are that neighboring cells have receptive fields from very similar parts of space, but have distinct responsibilities for those receptive fields, and that when neighboring cells do communicate, it is typically concentrated in delays of less than ±10 ms.

4 Technological Strategy

To have a hope of understanding local circuitry, we need to record simultaneously from multiple cells. Since we want to look at the local scale wiring constraint, we need to insure that our data originates from cells that are physically proximate. Multiple traditional extracellular electrodes, or even microwire electrodes, can be placed no closer than a few hundreds of microns apart with standard methods. Although portions of the typical cortical cell, such as the dendritic arborization, can extend up to the millimeter scale, the signal source is at the soma; in the macaque, cortical cells are packed at a distance of tens of microns apart, requiring an approach which uses electrodes spaced at approximately the same distance. Traditional single electrodes are sometimes capable of recording from more than one unit simultaneously which led, in part, to the development of multiwire bundle electrodes which use extremely fine microwires in close-packed configurations [41, 60]. These techniques allow simultaneous recording from multiple neighboring cells even in the densely packed hippocampus. We sought to adapt this technique to the awake behaving primate preparation, and a large portion of the dissertation will be devoted to this technology.

5 The Thesis

The primary thrust of this dissertation is to advance the understanding of neural circuitry. A secondary thrust is to reduce the invasiveness of the experimental techniques in common use, which benefits both experimenter and subject. We worked very hard to improve the way we collect data to allow this.

Using the knowledge acquired from a visit to McNaughton's laboratory at the University of Arizona, we set out to use the same recording technology to simultaneously observe multiple neighboring neurons in parietal cortex. The task chosen was the memory saccade, wherein an animal is instructed to delay a glance to a target until well after the target is no longer visible. As mentioned above, this task, in various forms and with great success, has been used to tease out different aspects of responses to visual targets in the parietal cortex. We developed a delivery method for the electrodes, we built equipment, verified operation in one monkey, tested our algorithms in locusts, and then and made a large number of recordings from two additional monkeys.

The presentation will follow the two themes of technological development and scientific exploration, presented sequentially in parts. The two parts, representing the two halves of the CNS program at Caltech, Engineering and Biology, are intended to stand on their own, but form a richer whole. To paraphrase Richard Andersen's advice on how to write this introduction, we have developed new techniques to address new questions, and we will tell you about both.

5.1 Part One, Technical Development

The first part consists of five chapters, all to address the central technological thesis that we can, indeed, record simultaneously from multiple cells in awake behaving monkeys with tetrodes. Chapter 2 (the first technology chapter) is an introduction to neuroelectrophysiology, concentrating on the main recording technique used in this thesis, the tetrode. Chapter 3 details the tetrode development done for this dissertation that is specific to the primary preparation used, including the mechanics of inserting tetrodes. Chapter 4 describes a new device for chronic positioning of multiple tetrodes. Chapter 5 documents some of the instrumentation built to acquire our data. Finally, Chapter 6 shows some of the workings of our spike sorting algorithm.

5.2 Part Two, Scientific Results

The second part consists of five chapters, which all address the central scientific thesis that neighboring cells in LIP have similar receptive fields but dissimilar response types. The first of these, Chapter 7, presents work done in collaboration with Michael Wehr to verify our recording technique by simultaneously recording from neurons extracellularly with tetrodes and intracellularly with sharp glass electrodes. Chapter 9 then reviews the experimental paradigms used to collect awake behaving monkey data from two animals. Chapter 10 discusses our results based on coarse time-scale analyses of neural responses to the task, and includes observations on similarities for response characteristics of neighboring neurons. Chapter 11 extends these results by examining fine time scale interactions between cells. Finally, Chapter 12 gives very preliminary results on multiple-tetrode recordings and looks to the future.

Everything is science fiction
— The Cars (I'm in Touch with Your World, *The Cars*, 1978)

Back off man. I'm a scientist.
— Dr. Peter Venkman (Ghostbusters, 1984)

Chapter Two: The Tetrode Technique

1 Introduction

For decades, the primary tool of the neurophysiologist has been the recording electrode, a device which transduces the electrical deviations of neuronal membrane voltage, allowing the experimenter to monitor a small handful of cells. Historically, this technique has been implemented primarily with a needle-like metal electrode, insulated along the shank up to the fine exposed tip; such electrodes are carefully brought in proximity to a neuron, generally any neuron within a few millimeter-diameter targeted volume, so that the signals from that neuron are large in comparison to signals from others and the background noise.

More recently, and originally to assist in isolating signals from neurons in the hippocampus, a brain area with high cell packing density, a technique using multiple electrodes in close-packed configuration either in pairs [41] or groups of four [60] was developed. These electrodes, under the names of *stereotrodes* and *tetrodes* respectively[1] allowed the simultaneous monitoring of multiple cells, a boon to the electrophysiologist. It is this technique, the use of electrodes in bundle form, that is the technological basis of this thesis, as will be described in detail below.

This chapter will introduce the technique and provide characteristics which should be applicable to many preparations, although some will be particular to the preparations used in the remainder of the document. The central five sections of this chapter will review, in turn, the idea of spheres of sensitivity, the impedances of tetrodes and traditional electrodes, the noise levels of tetrodes, the signal levels, and finally, the efficacy of tetrodes by way of the number of distinguishable cells per recording.

2 Spheres of Sensitivity

When an electrode is placed into active neural tissue, it is most sensitive to neurons directly against the tip, and progressively less sensitive to neurons further away. The sensitivity-versus-distance profile is often conceptually approximated as a thresholding function, and while there is substantial evidence that this is not true, we will not be examining the exact profile shape in the present paper. If the threshold approximation is taken in three dimensions, the resulting spheroid then circumscribes

---

[1] Terms such as *diode*, *triode*, *tetrode*, *pentode* and the like had been long-established within electrical engineering to describe vacuum tube devices with varying number of electrodes. The neuroscience community has adopted the ill-chosen terms *stereotrode* and *tetrode* to describe multi-wire bundle electrodes, mistakenly selecting the vernacular *-trode* as a suffix, rather than the preferable *-ode* (from the Greek, meaning *way* or *path*). For purely historical reasons, we will continue to use the established neuroscience terms for these electrodes.

the extent of detectable neuronal positions: neurons within the solid are electrically isolatable, those outside are not. For a more substantive examination of these issues, including asymmetries due to electrode configuration, see work by Kewley [31, 28]. Neurons closer to the center of the sensitivity sphere will have larger signatures than those farther away, and by simple counting arguments, we can deduce that there will be some characteristic distance, depending upon detection and identification algorithms, corresponding to the edge of the spheroid and beyond which individual neural signals become indistinguishable and are better modeled as noise. These two somewhat arbitrary regimes, signal and noise, underly the large body of electrophysiology, and form the basis for more advanced spike sorting algorithms, including the one developed as part of this project [67].

When multiple electrodes are positioned in neural tissue, we can separate the space of possible placement configurations into two classes: those where the sensitivity spheres of each electrode do not overlap, and those where they do. When the inter-electrode distance is sufficiently large such that the sensitivity spheres are non-overlapping, the difference between one electrode and many is insubstantial in terms of the information that can be extracted at each electrode. While there is a significant increase in total information obtainable by combining sorted spike trains with increasing numbers of such electrodes, additional electrodes cannot guide the extraction of additional spike trains unless our assumption of sensitivity spheres is invalid. Alternately put, the methods used to sort spike trains for a single electrode are equally applicable to multiple electrodes with non-overlapping sensitivity spheres, as they can be treated independently.

When the sensitivity spheres overlap, each electrode records slightly different versions of the overlapped electrical landscape, and the multiple views can be used to extract more information about the local activity than would be otherwise available. While single-electrode techniques can still be used with overlapped spheres, there is much more information to be gained by taking the signals in combination. An informal analogy can be drawn between these two cases, non-overlapping and overlapping spheres, and monophonic and stereo- or quadraphonic audio recordings: while a monophonic recording provides capable representation, the addition of spatial information available in multiphonic recordings provides substantially increased detail. Given our research goal of understanding local circuitry, it seems apparent that multiple electrodes with overlapping sensitivity spheres is the technique of choice, and we might therefore ask, how many electrodes should we be using?

In the simplified approximation of a noiseless point electrode listening to point sources in a homogeneous, isotropic, linear, passive, noiseless three-dimensional medium, is it necessary and sufficient to have four electrodes in non-planar configuration to uniquely identify the position of any arbitrary configuration of sources;[2] however, most of these assumptions are invalid for electrodes in

---

[2] A condensed version of the argument runs as follows. First, assume all neurons generate identical signals, all electrode tips are point transducers with identical characteristics, and signals decrease in a predictable way with increasing distance between source and transducer independent of direction. For a single electrode, all neurons positioned along

neural tissue. Real electrodes are noisy and have substantial extent; neurons are not point sources; and while there is some evidence that the medium is linear (see Wehr, Sahani, and Pezaris [72], and Chapter 7), the medium is inhomogeneous, non-isotropic, active, and noisy. While these non-idealities mean that using four electrodes *per se* is not as theoretically justified as we might like, for historical and practical reasons this electrode configuration formed the basis for our work. The preparation's inherent noise, combined electrode and neural, provides the background against which neurons cannot be distinguished once beyond some distance from the electrode, justifying the conceptual approximation of the sensitivity profile as a thresholding function. Within the sphere of detectability, we can expect some falloff of signal with distance, and therefore can use the characteristic amplitude and spike shape across the multiple channels to help identify signals originating from each neuron.

To accurately answer the question of numbers of electrodes, which we will not address in a satisfactory and thorough manner, we would need to consider not only the characteristics of the electrodes in a silent medium, specifically, the sensitivity falloff versus radial distance, but also the characteristics of the signal sources within the medium. If we assume that most signals are somatic (an approximation for vertebrate neurons), then the soma diameter and packing density become important. If the signals are not somatic, then we have a similar concern (density of sources) that is less easily measured. Additionally, we need to consider not only the inter-electrode distance, but the effective diameter at the electrode tip.

2.1 Beyond Four Electrodes

While traditionally the number of wires used in bundle electrodes is two or four, there is no inherent reason other numbers of electrodes should not prove useful. The limitation is dependent upon the wire diameter, the physical tissue damage caused by the bundle diameter, and the electrode packing density relative to the neuron packing density, as just mentioned. There has been work in both vertebrate and invertebrate preparations with differing numbers of wires in bundle form, and the performance varies substantially with the preparation. For example, Faller and Luttges [16] report over 20 isolatable cells per stereotrode in the dragonfly mesothoracic ganglia; McNaughton, O'Keefe and Barnes, in the first stereotrode publication [41] report up to 5 cells per stereotrode in the rat hippocampus; Recce and O'Keefe, in the first tetrode publication [60] report up to 10 cells

---

a sphere of radius r centered at the electrode tip will be indistinguishable, for any positive r. Adding a second electrode tip reduces the surfaces of confusion to circles centered on the line between the electrode tips, laying in the plane perpendicular to and bisecting that line. Adding a third electrode, the surfaces of confusion reduce to pairs of points (x, y, z) and their reflection through the plane defined by the electrode tips for any (x, y, z) not on that plane. Adding a fourth non-coplanar electrode establishes constraints that uniquely specify a neuron's position given the relative intensities amongst the four tips.

This argument for four electrodes relies on substantial and restrictive assumptions. Relaxing only two, namely, the assumptions of identical signals and identical transducers, complicates the argument immeasurably as neuron i at distance r_i may now appear the same as neuron j at distance r_j where $r_i \neq r_j$. Frankly, it's a wonder these things work at all.

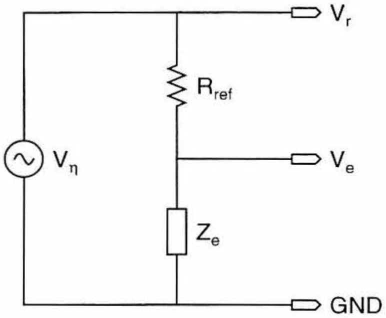

Figure 1: Simple Test Jig Schematic

per tetrode also in the rat hippocampus, and further report instances of improved tetrode versus stereotrode separability; Gray and colleagues [23] report 5 cells per tetrode in the cat preparation.

There are researchers who are planning larger multielectrode bundles, with as many as 10 or 20 wires in the primate preparation, and there are systems which may be purchased which include heptodes (seven electrodes, *e.g.*, Thomas Recording). At some point, care must be taken because the aggregate diameter of large numbers of fine wires will cause unacceptable damage to the tissue.

3 Impedance

When exploring the characteristics of a novel recording technology, the first question that should be asked is how well it transduces signals. In our case, the easiest way to do this is to measure the impedance of the electrodes over frequency to insure that transduction through tip capacitance dominates signal acquisition as expected. While we could also measure changes in the impedance for varying signal strengths, a second important issue, addressing that question was beyond the scope of our efforts; we limited our measurements to signal strengths that were approximately equivalent to *in situ* values.

We computed the impedance of electrodes by measuring the drawn current for a known applied voltage. This was done using a simple mechanism which applies a wideband noise source V_η to a voltage divider where one of the divider legs is a precision resistor R_{ref} of known value, and the other is the electrode under test (EUT, in deference to electrical engineering terminology of *device under test*, or DUT), as shown in Figure 1. The voltage across each leg of the divider is measured: the voltage V_r across the reference resistor gives the current through the test jig, which can be combined with the voltage V_e across the EUT to determine the unknown impedance Z_e.[3] These calculations are performed in the frequency domain, and used to generate a log-log plot of electrode impedance

---

[3] Voltage divided by current versus frequency yields impedance, or in this case, $Z_e(s) = V_e(s)/(V_r(s)/R_{\text{ref}})$, where s is complex frequency. In computing this we assume linearity and time-invariance to take advantage of the superposition principle and apply signal at all frequencies simultaneously. While it would be more accurate to measure the impedance at only one frequency at a time, allowing measurements of non-linear effects, with the equipment at hand, it was easier to use a broadband source and the technique described in the main text.

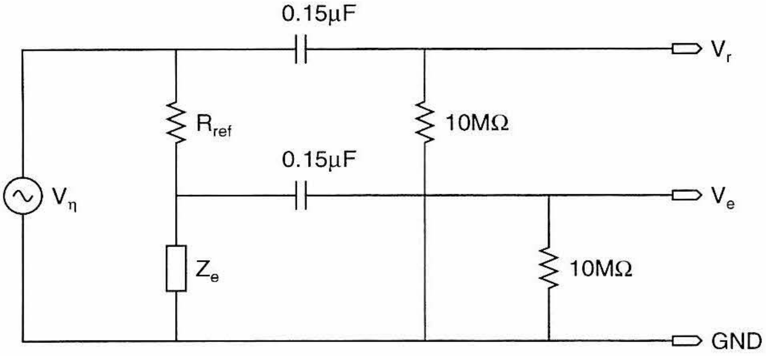

Figure 2: More Complete Test Jig Schematic

versus frequency. As long as the noise source bandwidth covers the frequency range of interest, the exact shape of its spectrum will factor out of the calculations. Similarly any non-zero finite source impedance may be safely ignored as long as it does not serve to suppress the effective applied signal below measurement thresholds.

However, finite input impedances of the recording apparatus cannot be ignored, as they will interact with the circuitry and serve to bias our measurements. Therefore, to accurately interpret the measured voltages, the test jig schematic must be augmented to include the input networks of the measuring amplifiers. This addition is shown in Figure 2 where we assume the input impedances of the operational amplifiers are sufficiently large as to represent an insignificant load, and therefore our recordings of V_r and V_e will accurately reflect the values at these nodes. The manufacturer's specification on the LMC6082AIM, the amplifier in question, is 10 TΩ for input resistance, and while not specified, typical input capacitances are in the range of a few picofarads, supporting the validity of this last assumption.

In full detail, our recording methods included a jig to hold the electrodes, a noise source, and a recording apparatus. The electrodes were placed, in turn, in the jig which held them a fixed distance above the bottom of a beaker. The jig held connectors, a 100 kΩ reference resistor R_{ref} (actual value 99.98 kΩ), and a bare platinum wire which ran into the bottom of the beaker and served as the ground node. A standard noise source (WG1, Tucker-Davis Technologies, Gainesville, Florida) was set to produce 1.0 $\mathrm{V_{p\text{-}p}}$ of uniform noise, which was fed to a precision attenuator (TDT's PA4) and brought down by a factor of 60 dB. This was delivered to the jig and applied as V_η shown in Figure 2. A multi-channel amplifier with input networks matching the figure buffered and amplified the signals $V_r(t)$ and $V_e(t)$ which were then fed to an apparatus to digitally capture them; this apparatus is detailed in subsequent chapters of this dissertation. The jig was placed in a Faraday cage connected to signal ground, and mechanically isolated from the room by a stack of foam and packing bubbles which rested on a heavy pedestal. To make a measurement, the beaker was slowly filled with saline while V_e was monitored, until a change was detected in signal level indicating the tip was just submersed. Recordings were made for 1000 s, digitized at 50 kHz with a 10 kHz anti-alias cutoff.

Additionally, to insure any asymmetries between the two recording channels did not adversely affect the measurements, a two-position switch was included in the jig to allow reliable and repeatable swapping of channels between the two nodes V_r and V_e. While some differences were found when interchanging channels, they were insubstantial compared to the electrode impedance, and restricted to the highest frequencies. It should be understood, however, that differences between the two recording channels, specifically in actual gains, frequency responses, and noise floors, set a baseline for the largest measurable impedance. For example, it was these effects that determined the measured gain for an open EUT (that is, no contact between electrode and saline bath).

Once the two voltages were measured, $Z_e(s)$, the impedance of the EUT versus frequency was computed as

$$Z_e(s) = V_e(s) \times \frac{R_1}{V_r(s) - V_e(s) \times \left(1 - \frac{R_1 C_1 s}{R_2 C_1 s + 1}\right)}, \tag{1}$$

where $R_1 = R_{\text{ref}}$, $R_2 = 10\ \text{M}\Omega$, $C_1 = 0.15\ \mu\text{F}$, and s, normally complex frequency, was restricted to the imaginary axis, $s = j\omega$. Were the input networks of the measuring amplifiers not taken into consideration, Formula 1 would reduce to $Z_e = V_e \times R_1/(V_r - V_e)$. The digitized streams were windowed in overlapping sections with a Hamming taper, Fourier transformed, and sections averaged to produce $V_r(s)$ and $V_e(s)$ above, from which $Z_e(s)$ was computed. The frequencies $\omega = 2\pi f$ for $f \in [10^0, 10^4]$ Hz were examined, for which the magnitude of the impedance was calculated.

The results for typical electrodes of varying composition are shown in Figure 3. Each electrode recording included a reference control recording of an open EUT to determine the maximum measurable impedance, made before the saline level was brought up to the electrode tip. The results were typically 10 MΩ or greater and reasonably independent of frequency. One of the control recordings is plotted in the figure, and can be seen to match the upper limit of resolvable impedance on all other traces.

Bearing in mind this upper limit, we see that all electrodes have a portion of their impedance spectrum which decreases, in most cases linearly on the log-log plot, with increasing frequency, indicative of an impedance dominated by a capacitive term. Nichrome tetrodes have the highest impedance, followed by stainless steel electrodes, then tungsten electrodes, and finally tungsten tetrodes. The static impedances measured at 1 kHz, respectively, 2 M, 1.5 M, 1.2 M, and 0.6 MΩ, directly correlate with the positions of the curves on the graph.

It should be noted that the impedance tester used to measure these 1 kHz values (IMP-1, Bak Electronics, Rockville, MD) has a DC-coupled zero-offset output stage which imposes a parasitic

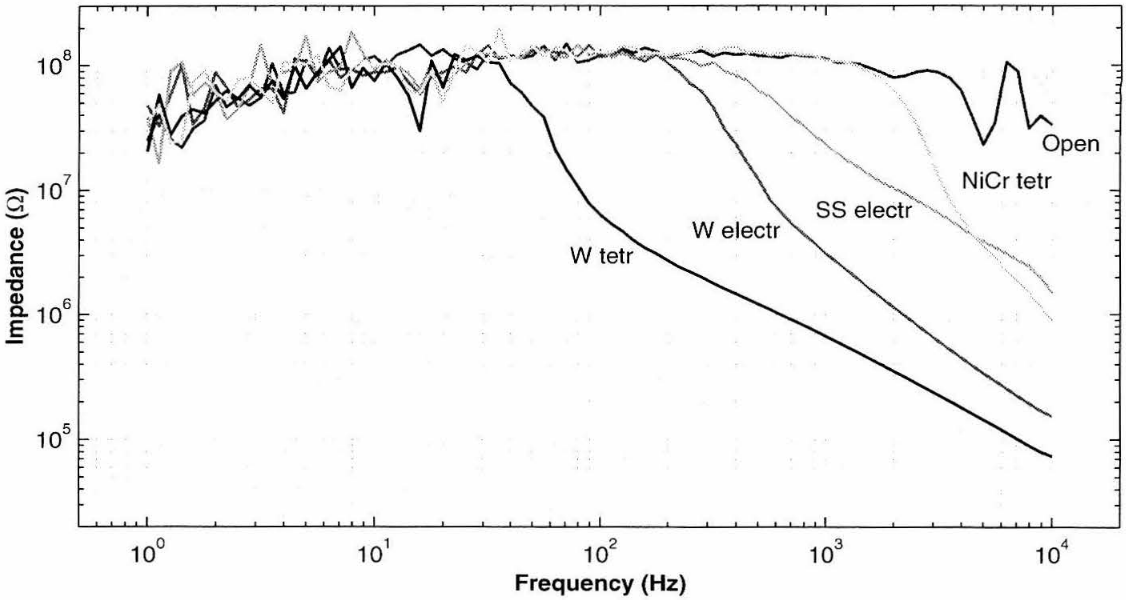

Figure 3: Electrode Impedances Over Frequency
The impedance characteristics of four electrode types (tungsten tetrode, tungsten electrode, stainless steel electrode, and nichrome tetrode) and an open circuit are seen versus measuring frequency. The open circuit trace sets the measurement limits, and the ripples at higher frequencies are due to differences between the two amplification channels used to capture the data (see main text). All four electrodes show a section of linear negative slope, indicative of a capacitance-dominated impedance.

current across the quiescent wet-cell voltage formed by the electrode tip, saline, and reference electrode, whereas the jig used here is AC coupled to avoid this problem. The result of the unintended DC current will be to electrochemically alter the tip and change the impedance over time; this is readily observed when using the IMP-1 as the displayed value changes over tens of seconds, decaying to lower values, often asymptoting to half the initial value. Although not presented, it was verified that the spectral values shown here did not change through the span of each recording, and across recordings for repeated measurements.

Interestingly, despite the common lore that higher impedance electrodes are better than lower impedance ones, we informally find that for tetrodes, the opposite is true: a lower impedance tetrode will produce better signal and better signal-to-noise ratios; however, our isolation methodology is substantially different from that used for a traditional electrode, as our placement is made to specifically avoid the dominance of one neuron's signals over those from others.

3.1 Tungsten Tetrode

For tungsten tetrodes, the primary electrode type used in this dissertation, we made a series of measurements to compare impedance due to tip capacitance against impedance due to capacitance through the insulation along the electrode wire. This was done by making impedance measurements for incremental levels of saline immersion in the same recording jig (the beaker was calibrated for height to volume, and measurements made at 3 mm increments from 0 to 12 mm). Although there were differences measured from depth to depth, they were not significant. The most direct interpretation is that tip capacitance dominates over shank capacitance in determining the electrode impedance, rather fortuitous for neuroelectric use, as it means signals from cells at the tip will be transduced and those along the insulated wire will be rejected.

3.2 Gold-Plated Nichrome

The nichrome impedance spectrum above showed an electrode of questionable utility which is corroborated by our experience: It is difficult at best to make recordings with plain nichrome tetrodes. To our knowledge, every researcher who uses nichrome tetrodes plates the tips with either gold or platinum. To verify that this serves to lower the impedance, we performed an experiment comparing the two. The impedance spectra are in Figure 4, where it can be seen that the impedance at 1 kHz has been lowered by two orders of magnitude from slightly over 1 MΩ down to 30 kΩ. While we did successfully make recordings with gold-plated nichrome tetrodes, the plating process proved difficult and unreliable; once we experimented with tungsten tetrodes, we did not again record with nichrome. This lowering of impedance qualitatively agrees with work done by Millar and Williams with plating carbon fiber microelectrodes with silver [42].

There are researchers (*e.g.*, [47]) who platinize tungsten tetrodes before use. This may have benefits for the chronic preparation in stemming imunoscavanging of the wire, but we find that tungsten tetrode signals are quite good without plating.

4 Noise

To properly distinguish neural signal from background noise, we must first fully characterize the noise. To this end, we made a series of recordings, in saline and *in vivo* neural tissue. Saline is often the medium of choice when making noise measurements. Our experiments suggest this is adequate to characterize the electrode for frequencies typically of interest to the electrophysiologist, but produces signal levels much lower than seen in living neural tissue.

Recordings were made from six tungsten tetrodes in saline using the same jig as used to measure impedances above, although the reference circuitry was not used, and no external voltages were applied. As before, recordings were digitized to 16 bits at 24 or 50 kHz with 10 kHz anti-alias filtering, and analyzed using in-house Matlab code to generate power spectra. Power spectra were computed using the same windowed methods as used for computing electrode impedance above.

The six power spectra appear in Figure 5, where it can be seen that above approximately 20 Hz, they are very similar, and exhibit the straight $1/f$ falloff one would expect from a capacitive

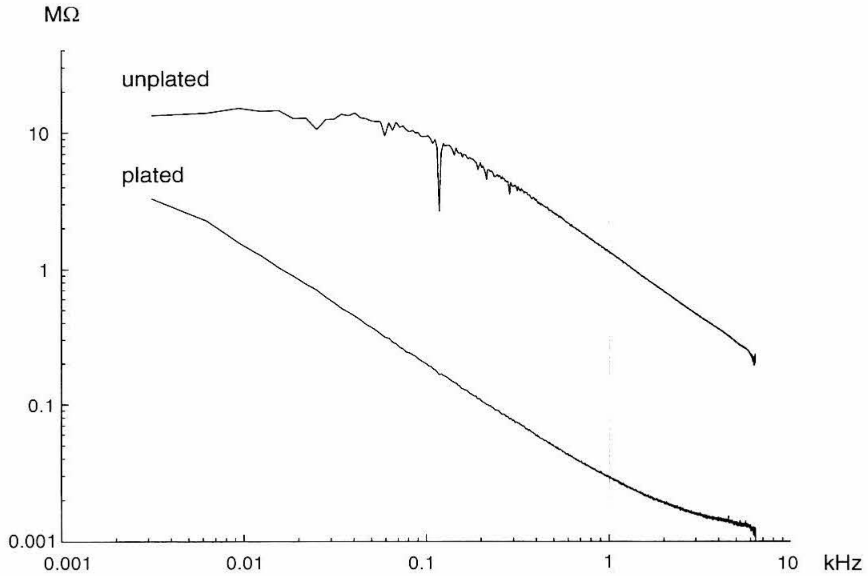

Figure 4: Effects of Gold Plating Nichrome Tetrodes
The upper trace is the impedance spectrum for a freshly cut nichrome tip in saline. The lower trace is for the same tip after plating in gold chloride ($AuCl_4$) at 2 μA for 4 seconds. The horizontal section of the upper trace is due to the limitations of our recording apparatus, as is the flattening in the lower trace at higher frequencies. Non-smooth features in the upper trace are due to power-line related interference (the largest dip is at 120 Hz).

load. Below 20 Hz, there is a high-energy process present in some recordings but not others. The multiple low-frequency peaks in this band suggest building vibration or similar mechanically-tuned effects. For the recordings without this low-frequency process, the $1/f$ falloff continues for nearly another decade downwards, below which the filtering effects of the recording apparatus start to come into play.

Two of these recordings were made directly following *in vivo* use of two tetrodes, allowing a comparison between neural and saline spectra on the same electrode. Figure 6 shows the traces, in mean form across the four tips, for the two cases. The conditions between the saline and neural recordings were matched as closely as possible; the animal was following the same behavioral task, the saline beaker was mechanically affixed to the primate chair, exposing it to similar, if not exactly identical, mechanical perturbation, and the same recording equipment was used with the same settings. The primary difference is that the saline recording has an order of magnitude less power than the neural recording. Both saline and neural recording spectra fall off linearly (in log-log space) above 10 Hz; however, where the saline recordings are quite straight, the neural recordings exhibit some deviations from straight-line descent, especially in the band 300–4000 Hz. We will examine these deviations in the next section.

The low-frequency energy in the saline recording in Figure 6A is similar the low-frequency

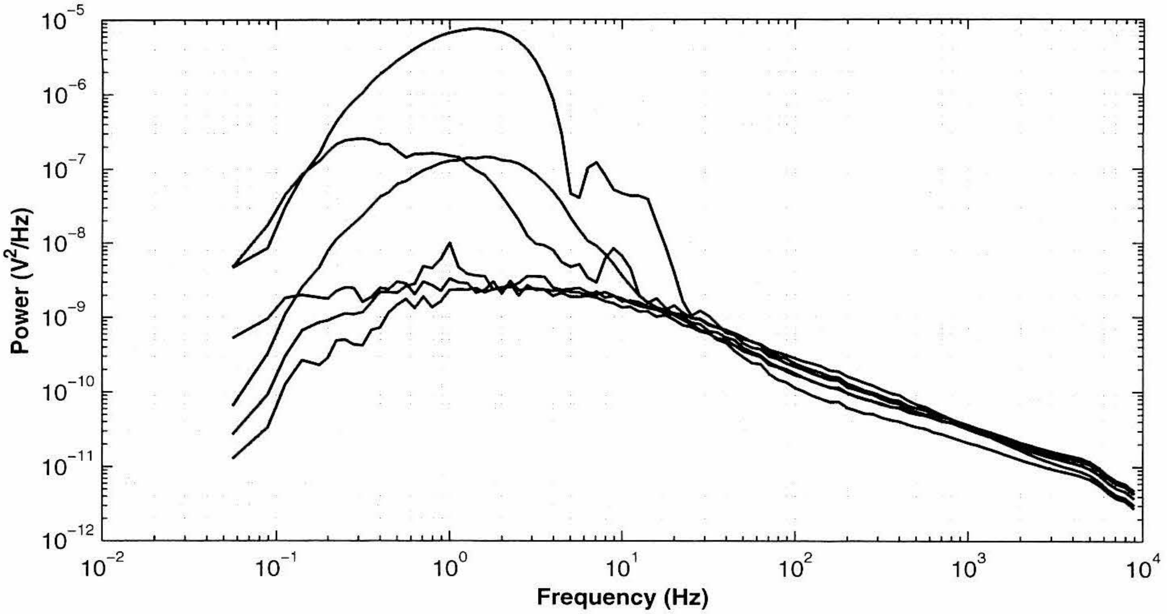

Figure 5: Tetrode Saline Noise
Spectra of six tetrode recordings made in saline. The four tetrode signals in a given recording were averaged together to generate each of the six traces here.

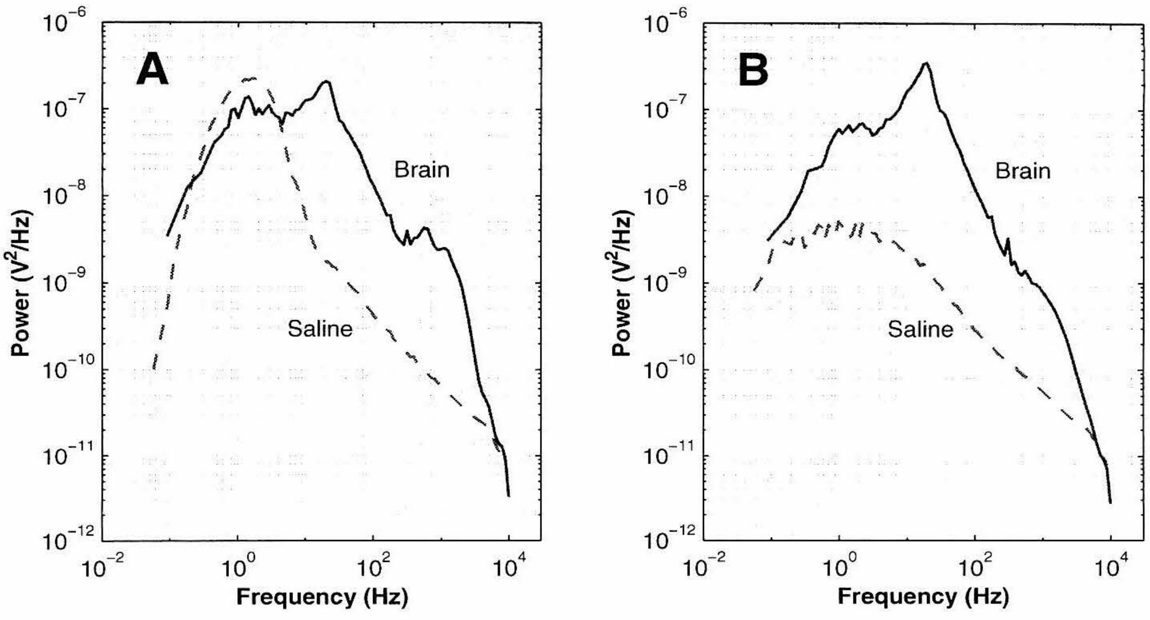

Figure 6: Tetrode Saline Noise and Neural Signal
Power spectra for neural recordings (solid lines; cmap4403, cmap4505) and directly subsequent saline recordings (dashed lines; saline05, saline06).

energy in the matching *in vivo* recording, but this is not true for the pair in 7B. We know by observation that the animal's movements during recordings (such as changes in posture, or sucking on the reward tube) are mechanically coupled to *in vivo* signals, thus we might speculate that the difference between the two pairs is one of quality of mechanical coupling to the animal chair during the saline recording. Specifically, the saline bath for the *in vitro* recording in Figure 7A was well-coupled, while the bath for 7B was not.

We conclude from these measurements that making saline recordings is useful for testing an experimental recording apparatus; however, as the power levels are approximately a factor of 10 lower, the voltage levels in saline will be a factor of three below neurogenic ones. This observation also suggests that neuroelectrical activity, as opposed to thermoelectric noise, should dominate neural recordings. This is all encouraging.

5 Signal

We will now address the hypothesis, established earlier in this chapter, that spiking energy can be detected above a relatively constant, or at least readily discernible, noise floor. In support of this hypothesis are analyses such as presented by Sahani, Pezaris, and Andersen [68] which examined the non-spiking activity in these recordings and found it to be Gaussian. We will extend this work by providing evidence that the excess high-frequency bump seen on the spectra presented thus far can be attributed to detectable spiking activity, first anecdotally through two examples, and then more formally by correlating a spiking index to spectral energy at 1 kHz against the number of detected spikes.

5.1 Example Recordings

We first examine two spectra, one from a recording without detectable spikes, and a second with many spikes, as seen in Figure 7. The spectra are nearly identical with the exception of an excess of energy in one in the range of 300–6000 Hz. Although these two recordings were made many days apart, the experimental conditions were similar (they were made from the same animal, in the same brain area, under the same task, and are of comparable length), suggesting that the difference seen in the spectra is related to the presence of spikes.

While suggestive, this example has many uncontrolled variables as the two recordings were made at different times. A more convincing argument would be made from recordings which show differences in spike amplitudes across the four channels, without differences in background activity. Two such recordings is shown in figures 8 and 9, where the background activity can be seen to be nearly identical across the four channels for the snippet of recording shown, while the amplitude of the spikes varies greatly. The spectra, on the right hand parts of the figures, show that at low frequencies, the four channels are nearly identical, while at higher frequencies (200–3000 Hz), the differences correlate strongly with the differences in spike amplitude. Every recording with spike amplitude differences shows similar differences in the size of this spectral bump, and no recording

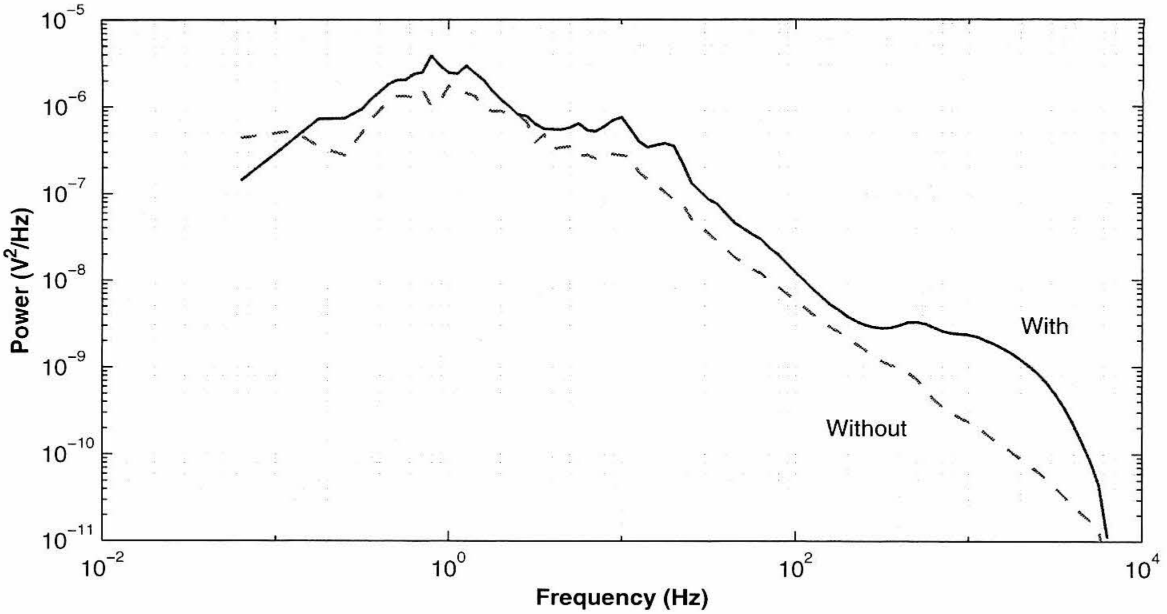

Figure 7: Example Tetrode Spectra With and Without Spikes
The two spectra in this diagram are from two monkey recordings (lmem2810, lmem5202) in the same animal, separated by many days. The first (solid line) has many well-isolated spikes, the second (dashed line) does not. Aside from an approximate factor of 2 difference in power, likely due to differences in electrode impedance, 0.6 and 1.3 MΩ at 1 kHz respectively, the largest difference is seen in the 300–5000 Hz band where the recording with spikes has substantially more power than the spikeless recording.

without differences in spike height shows localized amplitude differentiation in the spectrum. The presence of this effect in both vertebrate and invertebrate preparations suggests that deviations from straight line $1/f$ spectra are indeed neurogenic.

The locust example proves more compelling than the monkey example, as can be seen when directly comparing two other example spectra as in Figure 10. These spectra are means across the four channels for each preparation. The locust spectrum is much closer to a straight line with a smooth deviation 100–3000 Hz, while the monkey spectrum is less convincing, primarily because of low frequency deviations 4–30 Hz. While beyond the scope of this dissertation, we might speculate that this lower band is due to the local field signal, and the higher frequency deviations 300–3000 Hz are due to spiking activity. The examples in figures 8 and 9, where channels were plotted individually, support this, as we see the low frequency components are identical across the four channels, and the high frequency components vary according to the size of the spikes in the recording. But this remains anecdotal; we will examine the issue more rigorously in the next section.

5.2 Number of Spikes Influencing Spectral Shape

If we hypothesize that the total signal from a tetrode is a mixture of a background noise process and

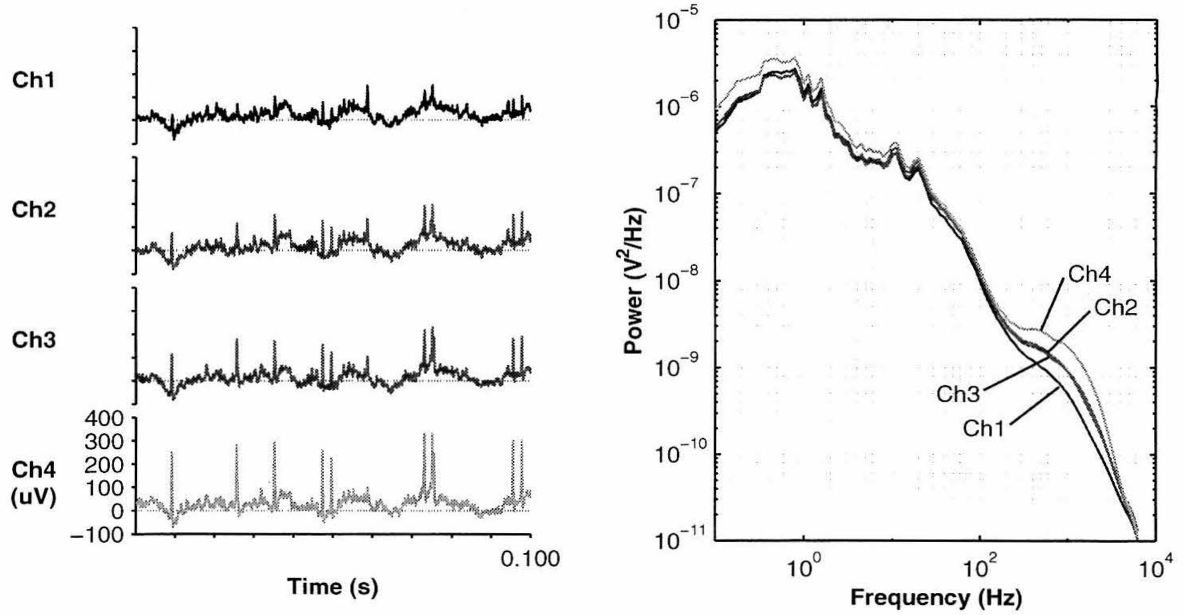

Figure 8: Example Monkey Tetrode Recording with Spectrum

The left panel shows a segment from an example tetrode recording (lmem4710) from one animal while the right panel shows the power spectrum for the recording. The spectral humps from 200 Hz to 4000 Hz show channel-to-channel differentiation which matches changes in spike size in the recordings on the left, suggesting a causal relation.

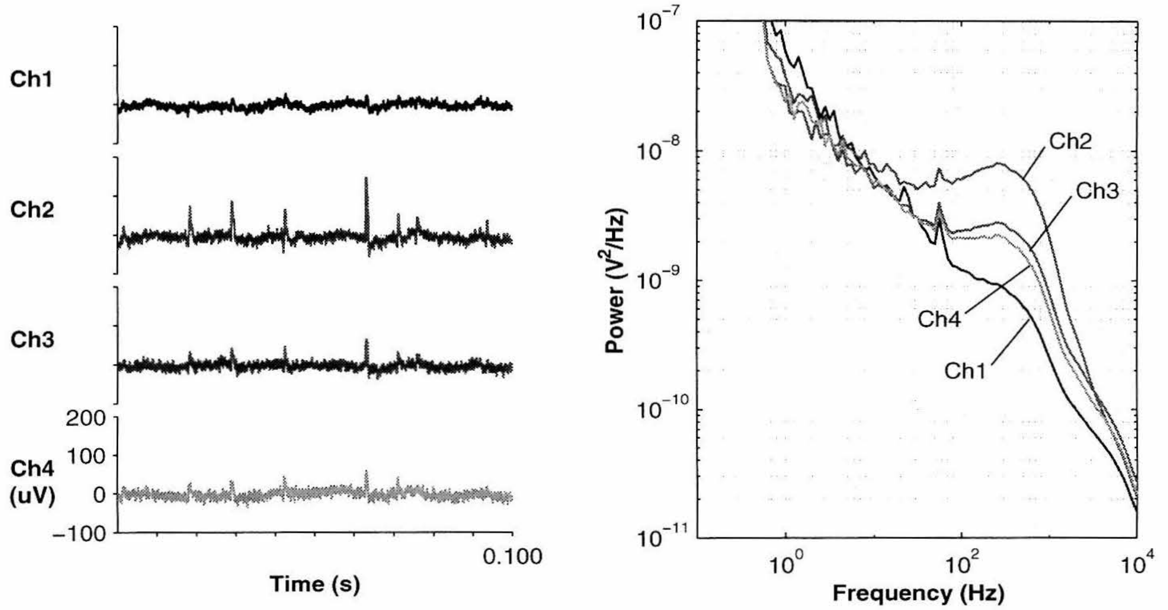

Figure 9: Example Locust Tetrode Recording with Spectrum

The same analysis as in the figure above was applied to signals from a locust recording. The left panel shows a segment from an example tetrode recording (pist02aA) from one preparation while the right panel shows the power spectrum for the recording.

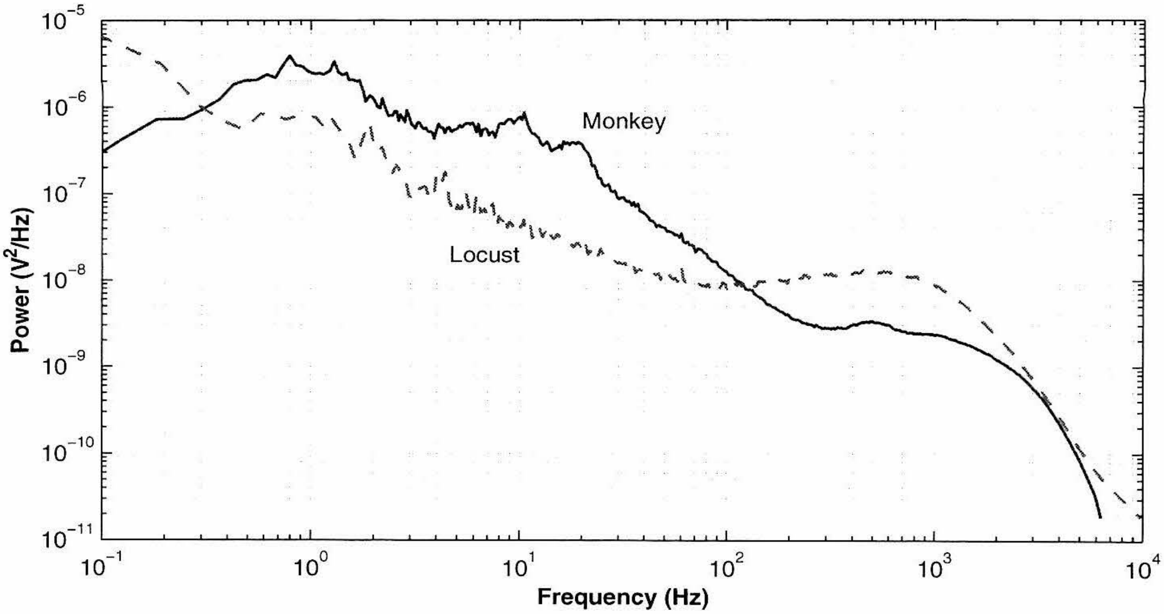

Figure 10: Monkey and Locust Spectral Signals
These two example spectra are for a recording in monkey neocortex (solid; lmem2810), and locust lobula (dashed; pist03aB), taken as the means across all four channels. They qualitatively show the same effect of a baseline linearly decreasing noise on top of which rides neural signal. The falloff in the monkey spectrum below 1 Hz is due in part to the low-frequency characteristics of the equipment used for that recording which differed from the equipment used for the locust recording.

a set of foreground spikes, that the spectrum of the background is close to straight line $1/f$, and the spectrum of spiking activity is concentrated in a band approximately 300–6000 Hz, then we would expect the fraction of a recording taken up by spikes should directly correlate with the deviation from $1/f$ spectra as the two sources mix. To properly model this, we would necessarily take account of the amplitude of each spike as well as the total number of spikes; however, for the present analysis, we will only examine the number of spikes normalized by recording length, assuming that each spike lasts a fixed amount of time, and that all spikes above a fixed threshold are the same size. By ignoring spike amplitudes, we should expect the observed correlation to be less than it would be otherwise, in particular, for evenly distributed data, we would expect to see a correlation coefficient of about 0.5, as the number of spikes and their amplitude should contribute equally.

To perform this analysis, we will take the suite of neural recordings which were collected for this dissertation, a portion of which is used later for spike train analysis, form their spectra, and correlate the mean firing rate (total number of spikes normalized by recording length) against the excess spectral power at 1000 Hz that lies above parabolas fitted to each spectrum. We use this power index instead of the unnormalized power to factor out differences in electrode impedance and other effects that will move the spectrum up and down without changing its shape; we fit with a

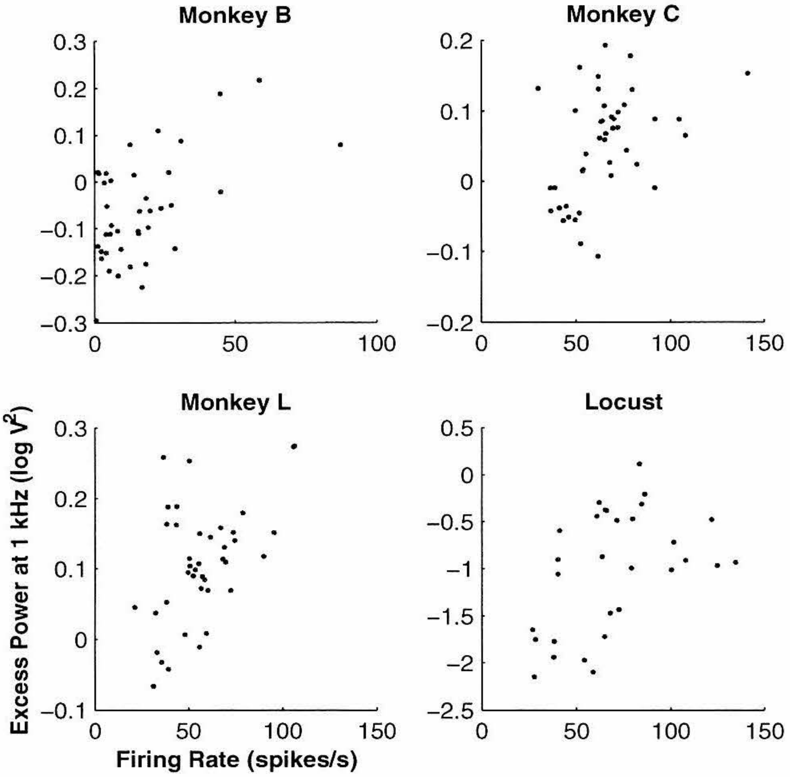

Figure 11: Power versus Firing Rate
Scattergrams of power at 1 kHz versus firing rate for collections of recordings in three monkeys and the group of locusts. Each distribution shows a significant correlation coefficient of about 0.4 (see main text).

parabola to extract only the size of the spectral bump. We use the firing rate as a proxy for the fraction of the recording which is spiking, as it is difficult to accurately measure the duration of each spike. Asymptotically, the firing rate should differ from the fraction of the recording only by the constant factor of the mean spike duration. The choice of 1000 Hz to examine the power is arbitrary, but based on the spectra seen thus far. If the mixing hypothesis is correct, we expect there to be a strong positive correlation between the firing rate and the power index. A scattergram of the results is shown in Figure 11. Recording methods are detailed in later chapters of this dissertation but broadly match the brief descriptions above. Spikes were detected by threshold excursion above 4 times the RMS value of the first 1 second of a recording for at least 2 consecutive samples on data digitally high-pass filtered at 300 Hz. For this analysis, data were filtered only for event detection, and left as-recorded for the spectral computations. Parabolas were fitted to each spectrum using the band from 0.1 Hz to 6400 Hz (the highest frequency common to all recordings), except for the locust data where instead of fitting a curve, normalization was done by subtracting the power in the band 1–100 Hz from each spectrum.

The results of this analysis show significant positive correlations for all three monkeys (B:

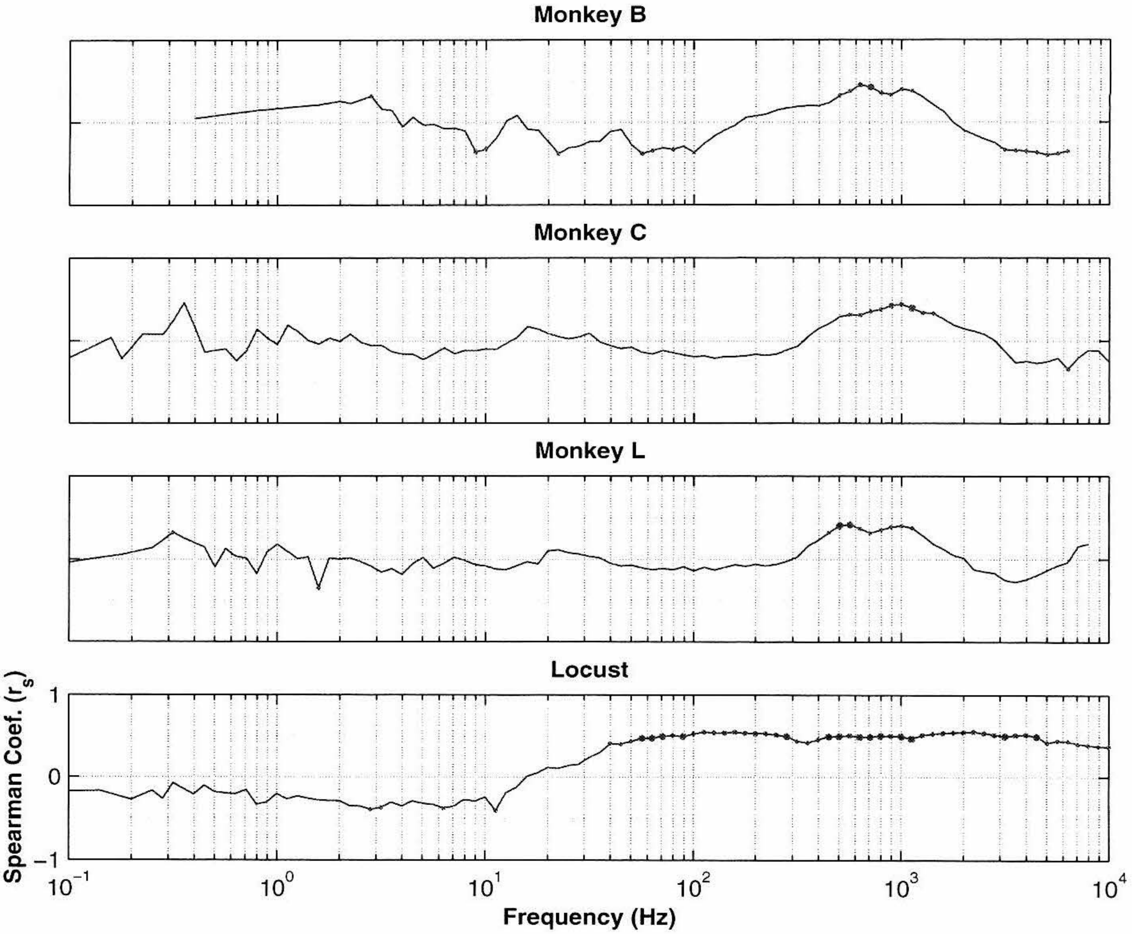

Figure 12: Correlation versus Spectral Frequency

Spearman rank coefficient (r_s, a non-parametric correlation coefficient) between firing rate and excess power ranging over different frequencies in the spectra for each preparation. The scattergrams in Figure 11 represent the data used to generate the points at 1 kHz on this figure. Dots represent significance values, small for $p < 0.05$, medium for $p < 0.01$, and large for $p < 0.005$.

$r_s = 0.40$, $n = 39$, $p < 0.05$; C: $r_s = 0.44$, $n = 43$, $p < 0.005$; L: $r_s = 0.41$, $n = 40$, $p < 0.01$) and the group of locusts ($r_s = 0.50$, $n = 30$, $p < 0.01$) as computed with the Spearman rank coefficient and shown in Figure 11. That is, nearly half of the variance in excess power is explained by variance in firing rate, as we would expect from the mixing hypothesis.

To verify the robustness of this finding, we performed the same computation at different frequencies. The coefficients plotted against frequency are shown in Figure 12, where all three monkey preparations show significant correlation in the band 500–1100 Hz, and the locust preparation shows significance from 40 Hz to the maximum recorded. The correlation for Monkey B also shows a significant negative effect at higher frequencies, above 3000 Hz, and spurious significances below 100 Hz, suggesting that the method chosen to normalize the spectra was not as appropriate for that preparation as for the other two monkeys. The recordings in Monkey B differ from those in Monkeys C and L in electrode material used (nichrome versus tungsten) and the level of effort made to isolate cells. This biases the Monkey B recordings towards lower firing rate sites with higher proportions of noise than recordings from the other two monkeys. Further, the recordings in Monkey B are substantially shorter than the others (*ca.* 100 s versus *ca.* 600 s) producing noisier spectral estimates. However, the three curves are qualitatively similar, showing little or erratic correlation below 300 Hz, a broad significant peak between 300 and 2000 Hz, and a dip above 3000 Hz.

A different normalization was used for the locust data because it was found that when fitting a line or parabola, the spectral bump deleteriously affected the approximation, reducing the observed effect. As the locust spectra are much simpler than the monkey spectra, and closer to a $1/f$ distribution, using a smaller portion of the spectra for a simpler fitting seems reasonable, and certainly the results are quite strong.

The breadth of the significance peaks and similarity across preparations suggests that we have indeed measured a robust phenomenon. If this is the case, then our mixing hypothesis would predict variations in spike-related power to be half from variations in firing rate, and half from spike amplitude. The measured shortfall from the theoretical $r_s = 0.5$ in monkeys could be due to two effects, firstly an inaccuracy in measuring excess energy, that is, an inaccuracy in modeling the background process when extracting the spike-related power, and secondly, experimental error in measuring the spectra. As the measured value for locusts nearly exactly matches the theoretical value, it is perhaps more likely that our background fitting could be improved for the monkey data.

It would be perhaps ill-advised to conclude from these significance plots that the frequencies of interest for spike sorting stop at 2000 Hz where the correlations on the monkey plots go to zero, for the exact shape of the curve is determined by the accuracy and appropriateness of the quadratic approximation to the spectra. The conclusion which should be drawn is limited to the significance of firing rate: for the monkey recordings presented, and the normalization methods chosen, firing rate has little predictive power over spectral shape above 2000 Hz. To fully understand the issues, we would have to also examine spike amplitude in these analyses. We would also be advised to examine the spectra of well-isolated spikes once extracted from the background noise. This will be seen in a later chapter.

5.3 Mean Spectra

To augment the findings thus far, we can also examine the mean spectra from the various animals used for the dissertation. All recordings which showed reasonable utility were gathered and the means taken for each monkey, and the group of locusts. Recordings with clear problems, such as clipping, no spiking signal, or those made differentially rather than single-ended as some of the earliest recordings, were not included. The spectrum of each recording was taken, and the means formed in logarithmic space. The results are shown in Figure 13, where all four traces can be seen to have an overall $1/f$-like structure. The mean locust spectrum has an increase in power at very low frequencies which can be attributed to recording artifacts due to respiration and changes in perfusion bath levels. Above 1 Hz, the mean locust spectrum shows a very nice straight line dropoff with an excess energy bump 50–3000 Hz, and all three monkey spectra can be approximated as straight lines descending downwards with two excursions, one in the range 10–50 Hz, and the second in the range 300–6000 Hz, although the evidence for such an excursion for Monkey B is weak at best.

The peak of the spike-related excursion, measured as the maximum deviation from an imagined $1/f$ baseline, occurs at a lower frequency for locusts, about 500 Hz, than for monkeys, about 1500 Hz. The upper limit on the spectra are all affected by limitations in our recording apparatus, and since the upper limits of the spiking features come very close to the edge of the spectra, we should remain cautious about determining an absolute upper end of the useful spectrum. The excursion extents, however, closely match the bandpass filter parameters traditionally used in single-unit isolation.

The recordings used for this figure from Monkey B were made with gold-plated nichrome tetrodes as part of our first experiments, while those from Monkeys L and C were made with tungsten tetrodes. This may, in addition to other technological advances made between preparations, explain the power differences between the mean spectra.

The low-frequency features in the tens of Hertz seen in the monkey spectra are, as previously mentioned, suspected to be from local field effects. Their absence from the locust may be due to the architecture of the area used for the recording, the optic lobula as detailed in a later chapter, having little synaptic structure. The rise in very low frequency energy in the locust spectrum may be due to respiration and other perfusion bath level effects. Although not shown here, it is also seen in simultaneously taken intracellular recordings (see Chapter 7).

6 Cell Count

Given analysis such as presented by Robinson [64] of cortical cell densities and electrode characteristics, and reasonable values for hearing radii, we expect to see approximately 8 cells per recording. Further, images such as Figure 14, and previous reports of tetrode efficacy [23], tend to support our expectation of a reasonably large number of cells per recording.

To verify this, we examined the set of monkey recordings used later in this dissertation for analysis relevant to a behavioral task, for the number of isolatable cells seen per recording site. These 47 recordings were sorted into spike trains using a statistical spike sorting algorithm [67] that will be described later, and the number of distinct cells, as identified by distinct cluster means, gathered. A range of 2 to 6 cells were found per recording, with a mode of 3 cells, and a mean of 3.0 ± 0.9 cells. Part of the original selection criteria for inclusion in the main study, including determining whether or not to take a recording at a given location, was to require a minimum of 2 cells per site, which

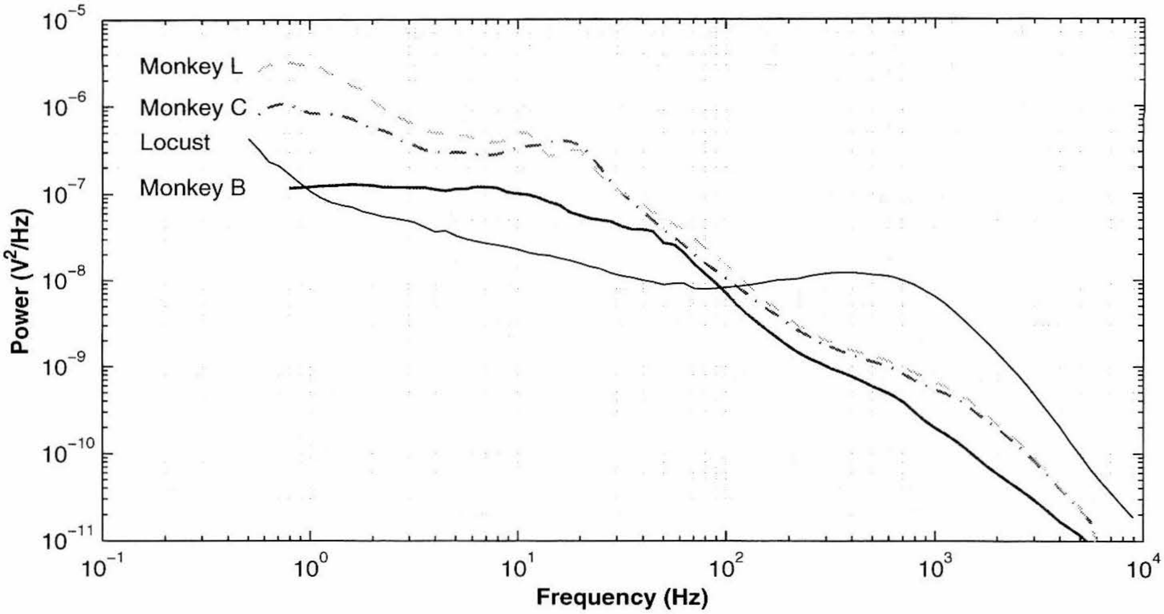

Figure 13: Mean Spectra
The mean spectra for recordings from three monkeys (B, n = 39; C, n = 43; L, n = 40) and a group of locusts (7 animals, n = 30 recordings). The progression of increased spectral power in the monkey recordings mirrors our increasing experience with the preparation and improvements in recording technique.

upwardly biases these figures; conversely, only cells which were cleanly isolated, as evidenced by a clear refractory period, are included in the counts, applying a downward bias. A histogram of the percent of sites with a given cell count is shown in Figure 15 where the results from Monkeys C and L have been pooled.[4] While there is not a significant difference between the distributions from the two animals (2-tailed t-test; individual distributions not shown), we do subjectively find that the number of cells per recording site has slowly increased over time, as we gain additional experience with the technique.

The average of 3–4 cells per site is substantially smaller than our prediction above based on Robinson's work. These observations, coupled with results from Tank and Kleinfeld [70] showing an *in vitro* hearing distance of 5 μm, and Buzsáki and Kandel [12] showing dendritic spike propagations detectable up to 400 μm from the soma, suggest that the oft-mentioned sensitivity radius of 65–100 μm is highly overstated. That larger value is anecdotally based on advancing electrodes through tissue and determining the positional extent through which individual cells are identifiable. We propose an alternate hypothesis to reconcile the apparently contradictory figures, namely, that electrodes have very small radii of sensitivity, explaining the apparently small number of observed cells per recording, and it is the somatic, dendritic, and axonal extent of each cell, that is, the

---

[4] Monkey B and locust results have not been included in the figure because they have not been computed. We would expect them to be shifted somewhat lower due to relative inexperience with the then-novel preparations.

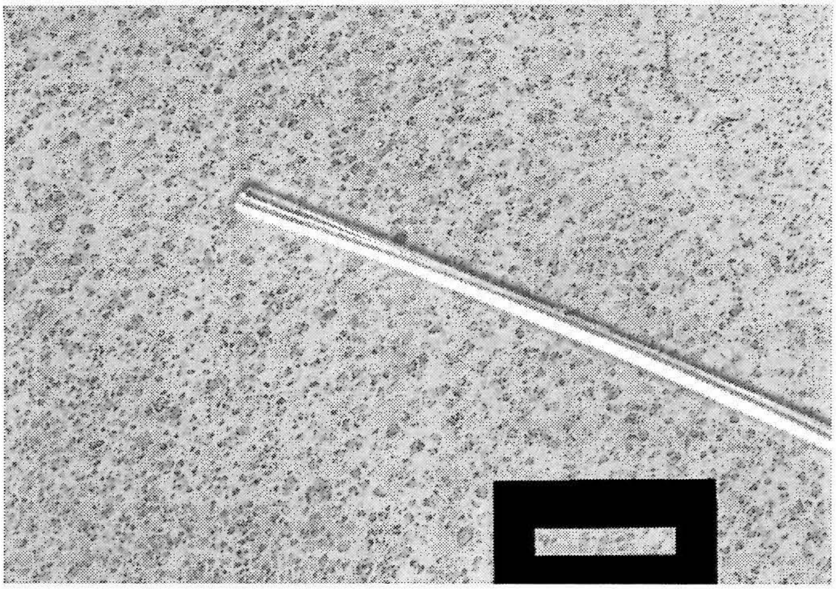

Figure 14: Tetrode Against Brain Slice
This is a dual-exposure image of a tetrode against a histologically-stained slice of tissue from the first tetrode recordings in monkey brain. The pair were first transilluminated to expose for the slide, and then epi-illuminated to expose for the tetrode. Nissl staining makes the neuronal cell bodies appear in this 50 μm thick slice as darker spots. The scale bar is 100 μm long.

non-point morphology of the sources, coupled with tissue pulling along with the electrode during advancement, more colloquially known as *brain drag*, which creates the substantially larger apparent hearing radius. Indeed, when passing through other parts of the brain, such as white matter, we regularly observe spikes, attributed to projection fibers, which appear and disappear in distances much smaller than the equivalent in gray matter. The morphological differences between projection fibers and soma, when met with a tetrode, we argue, provides an explanation for the apparent difference in hearing radius in this anecdotal example.

7 Summary

This chapter introduced the technique of recording with tetrodes, the basis of which is overlapped spheres of sensitivity. To understand spheres of sensitivity, the notion of separable foreground and background processes, in the guise of a threshold-like sensitivity-versus-distance profile for electrode tips in brain, was introduced, and evidence provided in support. The reasons for overlapping rather than distinct hearing spheres were discussed, and an analogy drawn to multichannel (*i.e.*, quadraphonic) audio recordings.

The impedance characteristics of tetrodes were presented, in comparison to other electrode types and our material of choice, tungsten, was justified due to its superiority over unplated nichrome. The tip impedance was shown to be primarily capacitive in origin in the frequencies of interest, and

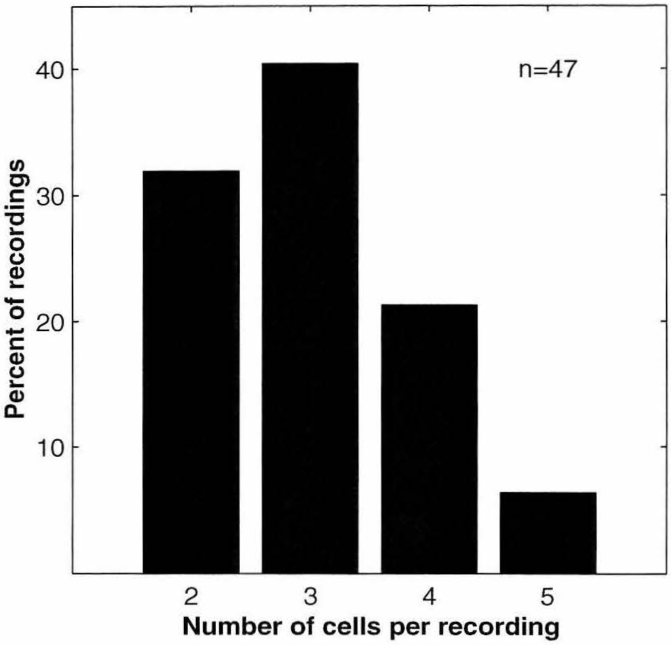

Figure 15: Cells per Site
Tetrode recordings in posterior parietal cortex currently yield an average of three cells per tip location in our lab, when selected for task response, multiplicity of cells, and ease of signal separation.

dominant over the shank capacitance. For the remainder of this document, tetrodes will be assumed to be made from tungsten wire, except where indicated otherwise. The noise from recording with tetrodes in saline was seen to be characteristic of a $1/f$ process, and substantially lower than the noise floor seen in live brain tissue.

The signal in brain tissue was shown to be concentrated in two bands, 10–50 Hz, and 200–5000 Hz, riding along a $1/f$-like background. The lower band was postulated to be from local field activity possibly contaminated with mechanically coupled artifact, while the higher band was shown to be linked to spiking activity both informally by amplitude and formally through firing rate. The spiking activity was seen to span a higher range of frequencies in monkeys than in locusts.

Finally, a profile of cell count per recording was presented for experiments done in two monkeys, along with arguments about the nature of and reported figures for the tetrode hearing radius. We postulated that the hearing radius, which with the morphology of the tissue determines the sphere of sensitivity, is actually quite small.

You'll never see eye to eye with the likes of us,
'cause we're too darn technical.
— THE HERBALIZER (Who's The Realest?, *Very Mercenary*, 1999)

Chapter Three: Tetrodes for Monkeys

1 Introduction

This chapter is based in large part on a technical report written in collaboration with Maneesh Sahani and Richard Andersen [56], and while some additional material has been added, specifically more detail on the *in situ* performance of tetrodes, the bulk of the material has been previously available.

In the preceding chapter, the method of recording with tetrodes was introduced and some of the basic electrode and signal characteristics presented. This chapter will discuss in detail the technology developed for delivering tetrodes into the awake behaving primate preparation, and the performance of the tetrode technique in our hands.

We have adapted the Recce-O'Keefe tetrode [60] for use in monkey cortex. Our design has been in use at this writing for five years in our laboratory, proving capable and robust. Details on construction of tetrodes, carrier tubes, and jigs are given in this chapter, including sources for materials. Example recordings show excellent isolation and histologically verified tracks are straight at up to 12 mm of tetrode extension, the maximum tested.

2 Background

Previously, tetrodes had been used in the mouse, rat, and cat preparations [60, 23]. On January 12, 1995, Maneesh Sahani and I made, to our knowledge, the first tetrode recordings in an awake primate preparation (technically, the recording only included 3 traces as the signal from the fourth channel was dead; the first true four-channel tetrode recording happened January 20, 1995, and is shown in Figure 16). This first animal was used for mechanical and electrical development and we recorded only spontaneous activity, as no behavioral controls were imposed. Despite this limitation, we were able to determine that we could, with certainty and reliability, detect multiple individual cells using this technique. The very first observation made during the first recording, after the euphoria of discovery and achievement dissipated slightly, was that spikes seemed to be riding along an underlying lower-frequency signal and correlated with it; the recording in Figure 16 shows hints of the effect. This observation was later confirmed by more formal examination of the coherence between spike trains and local field [48] in the suite of recordings that were made for part of this project.

An important realization regarding spike sorting was made during these early days: we started making differential recordings taken between tips, as opposed to single-ended between tips and

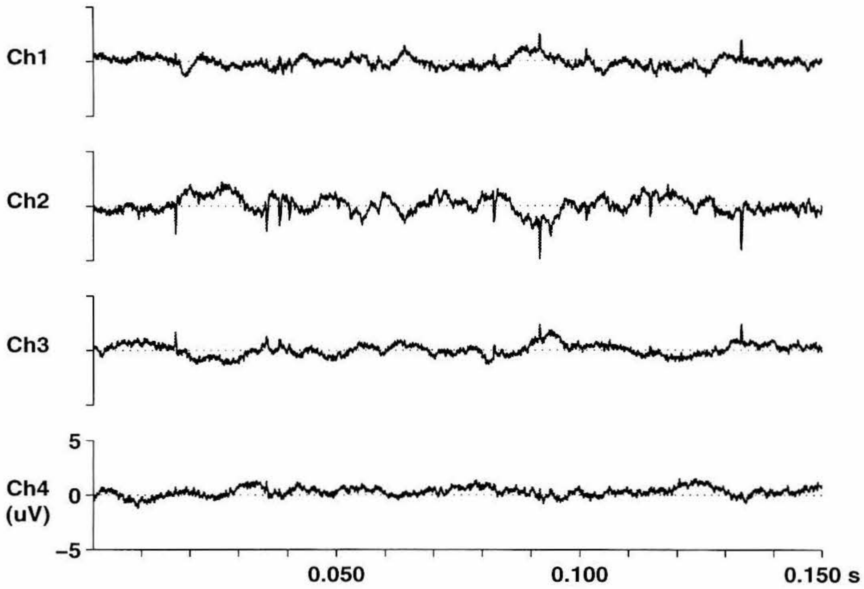

Figure 16: First Full Tetrode Recording
Made January 20, 1995, this segment is from the start of recording 11 for that day. While this was not the first tetrode recording made in our lab, the first attempt, a few days earlier, had only three channels operational out of the four. Spikes appear in both positive and negative directions as the recording was made differentially between adjacent electrode tips (eg, Ch1 shows the voltage between tip 1 and tip 2, Ch2 between tip 2 and tip 3, and so forth).

ground, to minimize interference from external sources, including mechanical artifacts. Regrettably, recording signals in this way induces an ambiguity as to neural spike polarity; whereas with single-ended recordings, each spike is observed as a positive deflection;[5] with differential recordings, some remain positive while others become negative. While we made attempts toward a general-purpose spike recognizer that would work for both positive and negative spikes, it proved difficult and the results were unreliable, thus the decision was made to compromise the advantages of differential recording for those of single-ended. Accordingly, the bulk of our recordings are made in single-ended fashion, and our spike sorting algorithms assume positive-going spikes.

Then, on April 22, 1996, we made the first tetrode recordings in an awake, behaving primate. A large number of recordings followed this, in a total of three animals, one for technical development (B), and two for scientific data collection (C and L). And, on May 21, 1999, we made the first dual-hemisphere, dual tetrode recordings in an awake, behaving primate, in a fourth animal (Y). This milestone was later revisited with other members of our laboratory almost one year later as dual-area tetrode recordings became routine.

---

[5] For conceptual ease, our recording apparatus includes an overall inversion so that voltages which are negative at the electrode tip appear as positive in our recordings. Thus normal extracellular action potentials correspond to upward going spikes in our graphs.

Adapting the tetrode technique to the awake behaving primate, and specifically for recording from the often difficult to find area LIP, meant creating a mechanism to reliably deliver the delicate tetrode bundle to a relatively deep structure. Additionally, as we elected to create a mechanism to make daily insertions, rather than maintain a chronically implanted electrode, we needed a way to protect the tetrode from handling damage while not in use.

Our eventual solution became known as the tetrode *carrier tube*, and is not unlike the traditional guide tube in that it is a sharpened stainless steel cannula which punctures the dura and other structures overlying the brain, and guides the tetrode's advance. However, the tetrode carrier is more complex than a traditional guide tube. A photograph is shown in Figure 17.

The carrier conceptually begins like a traditional guide tube, except that the electrode it carries is an order of magnitude smaller in diameter, being some 30 μm across, as opposed to 250 μm at the shank for traditional electrodes, thus, the walls of the upper end of the tetrode carrier tube (towards the right in Figure 17) are thicker than a normal guide tube. However, we sought to minimize tissue damage, so the carrier tapers down from a standard thickness (21 ga) where it mounts in the microdrive, to a fine diameter (32 ga) where it is inserted in the animal (towards the left in Figure 17). Additionally, the tetrode needs to be protected as it exits the guide tube at the upper end, and must be held in place against the hydraulic slave cylinder in the microdrive so that it can be actuated for penetration and withdrawal. Thus, a fine sliding cannula telescopes upwards out of the lower part of the carrier and is attached to a plate which is fixed to the slave cylinder during use. An internal stop within the lower part of the carrier tube prevents the telescoping portion from sliding out too far and separating from the lower part.

Using a carrier tube like this achieves the goals of full mechanical support for the tetrode bundle, protection from kinking and other damage during normal handing and preparation, the capability of targeting reasonably deep structures, and minimizing tissue damage during insertion. This final feature should not be overlooked, as, when properly sharpened, there is none of the traditional *pop* as when a guide tube penetrates the dura, and we often see no reaction from the animal subjects as the carriers are inserted.

3 Constraints Imposed by Animal Preparation

Previous reports describing tetrodes have been concentrated on rat [73] or cat [23] preparations. Both of these preparations featured chronic implantation of semi-moveable tetrodes using what is called a hyperdrive, and in both, while the electrodes are implanted chronically, the lifetime of the preparation is relatively limited. Two features of our awake monkey preparation which were in direct contrast to this previous work constrained our design. First, while a recording chamber is chronically implanted, the intracranial surface is exposed and electrodes are inserted and removed during each recording session. Thus, the dura mater must remain intact and is subject to toughening, in part as

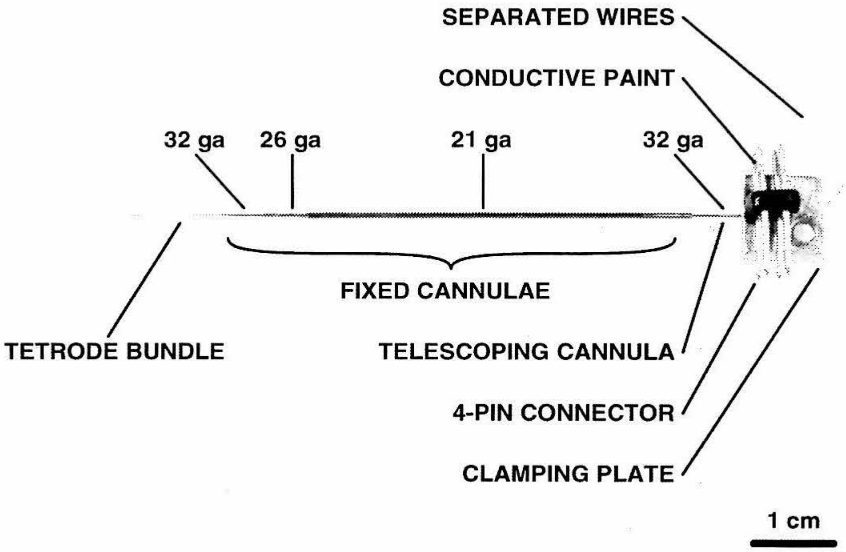

Figure 17: Tetrode Mk II Rev F, Photograph
An assembled and loaded tetrode carrier at approximately life-size. The tetrode bundle can be seen exiting the finest cannula on the left. The bundle runs through the nested tubes and exits at the other end of the carrier, behind the clamping plate. As the bundle exists, the wires are no longer twisted and fused as a bundle, and separate out to the four pins of the connector. The ends of each wire are attached both mechanically and electrically to the connector with conductive paint, after being chemically stripped of insulation. The nested cannulae on the left form the body of the carrier and are fixed by being glued to each other with epoxy. The telescoping cannula on the right is similarly fixed to the clamping plate, as is the four-pin connector, allowing the right-hand assembly to freely slide in and out of the body. Internal to the body, there are additional structures which prevent the telescoping section from completely coming out, or traveling in too far.

a result of the repeated penetrations. Second, the brain of the macaque is larger than those of either the rat or the cat and profoundly gyrated in comparison. Thus cortical areas can lie centimeters away from the exposed surface. A mechanism is therefore required to penetrate the dura and deliver the delicate tetrode wires to relatively deep neural structures, while minimizing tissue insult.

Our laboratory currently uses commercially available traditional electrodes and hydraulic microdrives (FHC, Brunswick, Maine), and we have designed our tetrodes to use the same equipment wherever possible. Specifically, our design will mount directly into an FHC microdrive, and is manipulated much like a traditional electrode would be. It could also be adapted to the Narishige or similar microdrive with little additional effort, with similar ease of handling.

Reiterating our goals in this chapter, before departing upon what will be an extensive description of our methodology, we seek to show that we can reliably deliver the fine tetrode bundle to structures below the cortical surface and make neural recordings with good signal and low noise. We will describe in detail the construction of tetrodes and carrier tubes, and then discuss the results of histological track verification, and typical signal to noise values for our recordings in monkey extrastriate visual cortex.

4 Methods

4.1 Introduction

Tetrodes consist of two parts, the bundle of four wires, often referred to as the tetrode itself, and a telescoping guide tube through which the tetrode bundle is advanced. Construction of the two will be described in the following sections, including details for the jigs used in assembly. We will first present construction of the special tetrode guide tubes, followed by construction of the machine to wind wire into tetrodes, and finally give instructions on winding and preparation of a tetrode.

4.2 Construction of Tetrode Carrier, Mk II Rev F

To protect the tetrode when not in use, to penetrate the dura, and to prevent buckling as the bundle is advanced, a stainless-steel assembly called a *carrier tube* is used. This *über*-guidetube, seen in figures 17 and 18, is constructed from readily available parts.

Carrier tubes are difficult to build as the parts are small and lengths and positions must be kept accurate to ±0.2 mm. A parts list is found in Table 1. The following detailed instructions assume that parts cut and machined to specification have been made or acquired. Exterior surfaces of lengths which will be glued (see Figure 18) should be roughed with 400 grit sandpaper. All lengths should be thoroughly cleaned (preferably ultrasonically) and dried.

Take a 10 mm length of 22.5° beveled 32 gauge cannula (C5) and slip the chamfered end into a 34 mm length of 26 gauge cannula (C4) to a depth of about 1 mm. Place a small amount of freshly mixed epoxy at the junction and slide into place until 5 mm of the finer cannula remains visible. Additional epoxy may be required as the two are put together; however, be sure to remove any excess. Call this assembly A1. Leave to dry.

Take a 5 mm length of 26 gauge cannula (C3) and similarly glue into place 18 mm from the blunt end of the 54 mm length of 32 gauge cannula (C6). Bring the smaller length almost into place before applying the epoxy to minimize any excess which must be removed. Call this assembly A2. Leave to dry.

Once A1 is dry, take a 49 mm length of 21 gauge cannula (C2) and glue into place 15 mm from the sharp end of A1. Call this new assembly A1′. Leave to dry.

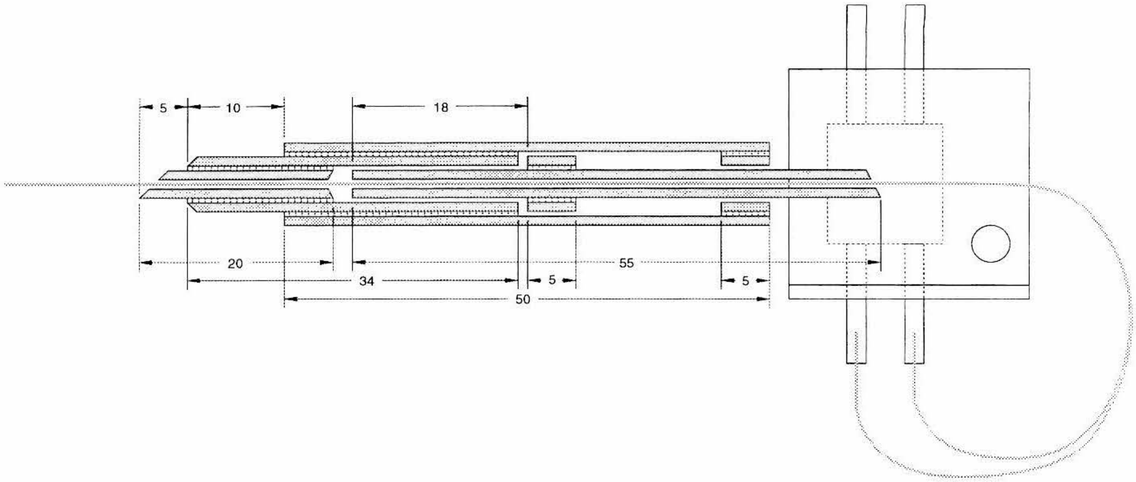

Figure 18: Tetrode Carrier Mk II Rev F, Section
The carrier tube punctures the dura directly, but is not advanced into the brain. The inner assembly is advanced into the carrier, forcing the tetrode bundle, which has been clamped to it when mounted in a microdrive, into the brain. All dimensions in millimeters. Hatching is glue. Heavy lines running through the center and curving out the right hand side are the tetrode bundle. The carrier is shown 1 mm from the closed position. Drawing not to scale (vertical dimensions exaggerated in portions). In this drawing, the tetrode has been rotated 180 degrees along its longitudinal axis relative to the photograph in Figure 17.

item	*quantity*	*description*	*source*	*part number*
C1	1	clamping plate	handmade	
C2	1	21 ga, 50 mm long[a]	Small Parts	O-HTX-21-12[e]
C3	2	26 ga, 5 mm long	Small Parts	O-HTX-26-12[e]
C4	1	26 ga, 34 mm long[a]	Small Parts	O-HTX-25-12[e]
C5	1	32 ga, 20 mm long[b]	Small Parts	O-HTX-32-12[e]
C6	1	32 ga, 55 mm long[c]	Small Parts	O-HTX-32-12[e]
C7	1	connector[d]	DigiKey	S2231-36
C8	misc	quick-setting epoxy[f]	Devcon	S-208
C9	misc	soft wax		
C10	misc	cellophane cooking wrap		

Table 1: Tetrode Carrier Parts List
[a] OD chamfer one end. [b] ID chamfer one end, 22.5° bevel cut other end. [c] 45° bevel cut one end. [d] Cut 4 pins (2×2) out of overall length. [e] Cannulae part numbers are for uncut 12-inch lengths; Small Parts will provide lengths cut and machined to specification at additional cost. [f] Trade name 5-Minute Epoxy.

Once A1′ is dry, insert A2 into A1′, beveled end outwards. Then, take a 5 mm length of 26 gauge cannula (C3) and slide over the free end of A2 and glue to the *outer* cannula of A1′ (the length of 21 gauge) so that the ends are flush. Leave to dry.

Once the epoxy has set, the carrier tube should telescope smoothly approximately 15 mm from stop to stop. The two parts should also rotate freely against each other. The remaining step is to glue the free end of A2 to the clamping plate. To prevent glue creeping into the hole during this process, seal it with a bit of soft wax. To align the carrier tube to the clamping plate, this last step is best done using a microdrive as a jig. Insert the carrier into the guide tube hole and fix in place by *lightly* tightening the microdrive set screw. Slide A2 until it extends 5 mm over the edge of the microdrive clamping block, rotating it so that the bevel is towards the block. Slide a bit of cellophane under the end of A2. Apply a *small* amount of glue to the 32 gauge cannula and screw the clamping plate down. Glue the 4-pin connector to the facing side of the clamping plate in a transverse orientation (it may be necessary to put a very small section of circuit board underneath the connector as a spacer) using quick-setting epoxy. Leave to dry.

At removal, carefully peel off the plastic sheet, insert a reaming wire from the penetrating end and work out the wax sealing the opposite end. Thoroughly clean and lubricate before use (see below).

Note that because the glue used in construction is non-conductive, the outermost section of the carrier (the 21 gauge cannula of A1) may not be in electrical contact with the innermost (the 32 gauge cannula of A1), and therefore may not be in contact with tissue when in use. It is advantageous to ground this cannula from a shielding standpoint, but it should not be used as a source of ground unless additional steps are taken, such as using a conductive epoxy for assembly or filling the recording chamber with saline during use. Our initial experience with conductive epoxies has been disappointing, and given their excessive cost, it seems prudent to use non-conductive glue and assume the lengths of cannulae are all isolated.

4.3 Carrier Tube Cleaning and Lubrication

Carrier tubes are cleaned in four steps by repeated reaming with music wire first wetted with acetone, then dry, then wetted with isporopyl, and again dry (this is all done with the same length of wire). They are then lubricated by reaming with wire wetted with silicone oil. This litany should be done before the first use, and every time the tetrode bundle is replaced. Oil-filling the carrier tube is important to increase construction yield, reduce the chances of electrode kinking during use, and prevent wicking of physiological fluid up the carrier.

4.4 Tetrode Twister

Tetrodes are built by winding fine wire on a motorized jig called a *tetrode twister*. Our twister was constructed out of readily-available inexpensive parts and provides sufficient flexibility for nearly any taste. The twister (see Figure 19) has two glass hooks, one suspended some 15 cm above the other, arranged on a chemistry stand. The upper hook is clamped to the stand, and therefore adjustable in height but normally remains fixed; the lower hook is attached to the output shaft of an electrical motor, and can be rotated about the vertical axis in either direction. Parts required for construction

item	*quantity*	*description*	*source*	*part number*
T1	1	chemistry stand		
T2	2	clamp (one machined)		
T3	1	132 RPM reversible AC motor	C&H Sales	ACGM9750
T4	1	DPDT center off 10 A switch	C&H Sales	SW9603
T5	1	three-pronged power cord	C&H Sales	PFC8300
T6	1	shaft coupling[a]	C&H Sales	
T7	1	5 cm × 5 cm blank PC board		
T8	2	glass hooks[b]		
T9	2	plastic tubing[c]		
T10	misc	wire, solder, electrical tape		

Table 2: Tetrode Twister Parts List
[a] *To fit output shaft of motor.* [b] *Bend from glass rod using a Bunsen burner.*
[c] *Sized to fit glass rod; 5–10 cm long.*

are listed in Table 2. Some machining of the lower clamp is required, and the switch must be glued to the motor along with a guard.

4.5 Twisting a Tetrode

Twisting a tetrode is a straightforward but delicate operation. It must be done with clean hands. First, place a 5 cm strip of adhesive tape on the motor from the rear, leaving the ends free. Draw a 60 cm length of tungsten wire and cut with fine surgical scissors (use an inverted glass funnel as a low-friction bearing for the spool of wire). Place one end of the wire against the adhesive of one end of the tape and press the tape against the motor shaft. Carefully wrap the wire from the lower hook to the upper hook, making two loops. Bring the wire just barely taught, and wind 4 or 5 times just above the lower hook before fixing it to the shaft with the other end of the tape. Engage the motor in the forward direction for approximately 65 turns. After winding, evenly pass a heat gun (FIT GUN 3, Alpha Wire; set to *high*) over the wire at a distance of 5–10 cm taking 5 seconds for each of two passes. Carefully cut the tetrode free from the bottom (it may help to hold the bundle just above the cut point with a pair of plastic-coated forceps, releasing the tension slowly after the cut). The tetrode will now appear as a single bundle of wire with a loop at one end; a properly twisted tetrode is dead straight. Once wound, take care not to bend or kink the tetrode bundle; a number 10 envelope makes a nice storage container.

4.6 Stripping

Electrical connections are made to the looped end of the bundle (as it comes from the twister). With a pair of very sharp scissors, cut the loops and trim the wires even. Place the tetrode on a pad of paper, and cover the final 3–5 mm of wire at the (now cut) loop ends with chemical strip (Strip-X, GC Electronics) for 20–30 minutes. Gently pull the wire out of the chemical strip, discard the upper sheet of paper and lay the tetrode bundle back on the pad. Take one cotton swab and

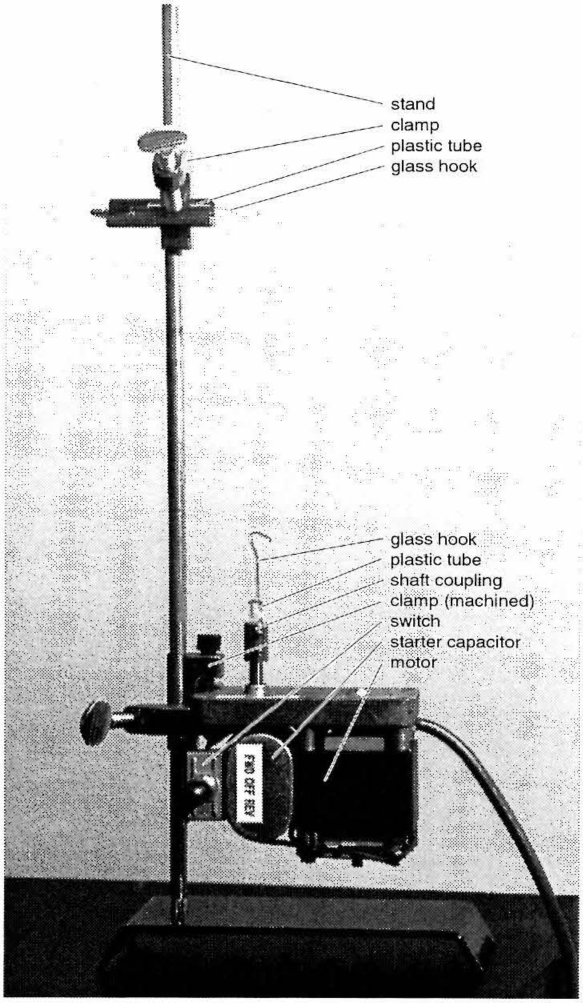

Figure 19: Tetrode Twister
This jig is used to twist tungsten wire into tetrodes. Wire is looped between the two glass hooks, held in place with adhesive tape, and then twisted when the motor is engaged. Forward and reverse directions are by convention only.

use it to hold the bundle against the pad (press the cotton bud quite hard against the wire about 1 cm below the stripped area); take another cotton swab, dip it in isopropyl alcohol and wipe any remaining chemical strip off the wire. Finally, check the bundle under a low-power microscope to insure the insulation has been properly removed.

4.7 Loading

Loading a tetrode bundle into the carrier should be done under low-power magnification (4× to 8×) by someone with steady hands. Caffeine use is discouraged. Grasp the recording end of the bundle with a pair of plastic-coated forceps and approach the clamping plate end of the carrier. Carefully,

Figure 20: Tetrode Tools
The set of hand tools used to create tetrode bundles and carriers. From the left, a pair of labelled sharp surgical scissors to cut tetrode wire, a pair of strong forceps for manipulating reaming wire and the like, a pair of blunted fine forceps for manipulating stripped tetrode wire, and a pair of fine forceps with Tygon tubing forced over the tips for manipulating insulated tetrode wire.

and repeatedly, attempt insertion until the wire has been threaded. If the wire gets kinked, trim the damaged part off before continuing further attempts. Once the bundle has been threaded, use the forceps to slide it into the carrier, a few millimeters at a time. Be careful not to bend or kink the wire during this process. Eventually, the bundle will encounter the lower end of the innermost cannula, and it may prove difficult to thread. Patience! Having ID-chamfered the innermost cannula (beveled the inner edge) will help. Try extending the carrier to different lengths, holding the bundle against the clamping plate with the forceps, and using a combined sliding and rotating motion on the lower part of the carrier to thread the wire, watching for bends where it exits the carrier at the top. It may become necessary to remove the bundle and trim the end if it proves impossible to thread through the last section. Once the recording tip appears at the sharp end of the carrier, pull the bundle through until 2–3 cm remain at the other end and proceed to the next step. Be careful not to bend or kink the wire.

4.8 Making Electrical Connection

The stripped wire is physically and electrically attached to the connector with conductive paint. The connector should first be cleaned by dipping the head of the carrier in isopropyl alcohol, and then dried. Once dry, each wire (previously stripped) is attached to the connector by placing a drop of conductive paint on a connector pin, and, again under magnification, one of the four wires brought into the paint and held in place until the paint dries. Clean uncoated blunted forceps should be

item	*quantity*	*description*	*source*	*part number*
T1	1	Dumont Pattern 1[a]	BRI	10-1400
T2	1	Dumont Pattern 5	BRI	10-1425
T3	1	Dumont Pattern 5[b]	BRI	10-1425
T4	1	micro dissecting scissors[c]	BRI	11-2030
T5	1	lamp with magnifier glass[d]	Luxo	17253BK
T6	1	sharp wire cutters[d]	Xcelite	MS54
T7	1	heat gun	Newark Electronics	93F1914
T8	misc	conductive paint[e]	GC Electronics	22-201
T9	misc	chemical strip[f]	GC Electronics	10-2602
T10	misc	tungsten wire	California Fine Wire	CFW211-0005-HML
T11	misc	acetone		
T12	misc	isopropyl alcohol		
T13	misc	reaming wire[g]		

Table 3: Tetrode Tools and Supplies
List of tools and supplies used to twist bundles, construct carriers, and to load bundles into carriers. [a] *Blunt tips slightly with a file or grinding wheel.* [b] *Cover tips with fine surgical tubing.* [c] *Label* **TETRODE WIRE ONLY** *and separate from other tools.* [d] *Available through DigiKey.* [e] *Trade name* Silver Print, *available through electronics distributors and retailers.* [f] *Trade name* Strip-X, *available through electronics distributors and retailers.* [g] *This should be music or spring wire no larger than 0.003 inches in diameter.*

used for this, and again, patience is necessary to hold the wire as the paint dries (10–15 seconds is usually sufficient to make a physical connection; full electrical connection takes several minutes). An articulated-arm incandescent lamp with conical shade makes an excellent low-temperature hood to speed drying. See Figure 19 for an assembled and loaded carrier.

4.9 Tools and Supplies

Recommended tools and supplies to construct tetrode bundles, carriers, loading and cleaning can be found in Table 3. Most should be available in biology and electronics stock rooms; however, sources for nearly all are listed at the end of the paper.

Reserve the fine surgical scissors for cutting tetrode wire only. Do not use them for any other purpose, as they must be extremely sharp to cleanly cut the fine tungsten. A pair of scissors lasts approximately one year in our laboratory; once they become too dull for this task, they remain quite useful and can be employed elsewhere.

4.10 Recording

Our experimental setup is reviewed in intimate detail in subsequent chapters, but will be summarized here. A hydraulic microdrive (Fred Haer Corp, Brunswick, Maine) is used to position the tetrodes. Tetrode signals are amplified by a custom four-channel headstage amplifier ($A = 100$) feeding a custom four-channel variable-gain preamplifier ($A = 1$ to 5000, nominally set to 200). The preamplifier feeds anti-alias filters ($f_c = 10$ kHz, Tucker-Davis Technologies, Gainsville, Florida) and four-channel instrumentation-grade A/D ($f_s = 20$ kHz, also TDT). Data are streamed to disk, eventually written to CD-ROM, and analyzed using in-house MATLAB code. Recorded streams are parsed into spike trains using spike sorting algorithms which have been described elsewhere [69, 68, 67], and are presented in a subsequent chapter.

4.11 Track Recovery

Evidence from informal testing with an artificial brain substitute (household gelatin, mixed to double strength) suggested that tetrodes, if adequately supported above the brain, would penetrate without deviation once in tissue without additional support. We therefore made two groups of penetrations in brain marked with electrolytic lesions, in two animals, separated by a number of years, to verify this. Lesions were made with tetrodes prior to sacrifice and perfusion with formaldehyde. In the first preparation, five penetrations were made with nichrome tetrodes, and two control penetrations with traditional tungsten electrodes. Lesions were made with spacings of 2 mm, 1 mm, and 0.5 mm, with most penetrations having two or four lesions at the largest spacings, and two penetrations having over 10 lesions at the finest spacing. Lesions were made at 8 μA for 10 seconds, passing current into the electrode and recovering at a previously implanted distant skull screw. Lesions were made within 1 week of sacrifice.

In the second preparation, seven penetrations were made with tungsten tetrodes: two penetrations had punctate lesions every 2 mm; three had lesions every 0.5 mm; two had constant lesioning current applied as electrodes were smoothly withdrawn from the deepest point. For punctate lesions, 8 μA of current was passed for 5 seconds; for continuous lesions, current was passed at 4 μA while withdrawing at approximately 1 mm/min. All penetrations were made to 10 mm of tetrode extension, starting at or just below the cortical surface, with lesions spanning the entire penetration. Lesions were made within 4 days of sacrifice.

5 Results

An example tetrode recording is shown in Figure 21. This recording is of particular interest because it shows one unit, a, detected on channel 3, a second unit, b, on channel 2, while both appear on channel 1 at approximately equal amplitude. Two additional units, c and d, were detected in this recording. Notice how similar three of the four waveforms are on channel 1, making it very difficult to disambiguate between them given only the information from that trace. For pedagogical reasons, this example was chosen because of the unusually large separation between channels; more commonly, the distinction between channels is not as extreme as here; however, the signal to noise ratio seen in this segment is not uncommon among our tetrode recordings.

As we continuously capture the electrode voltage, rather than the traditional approach of

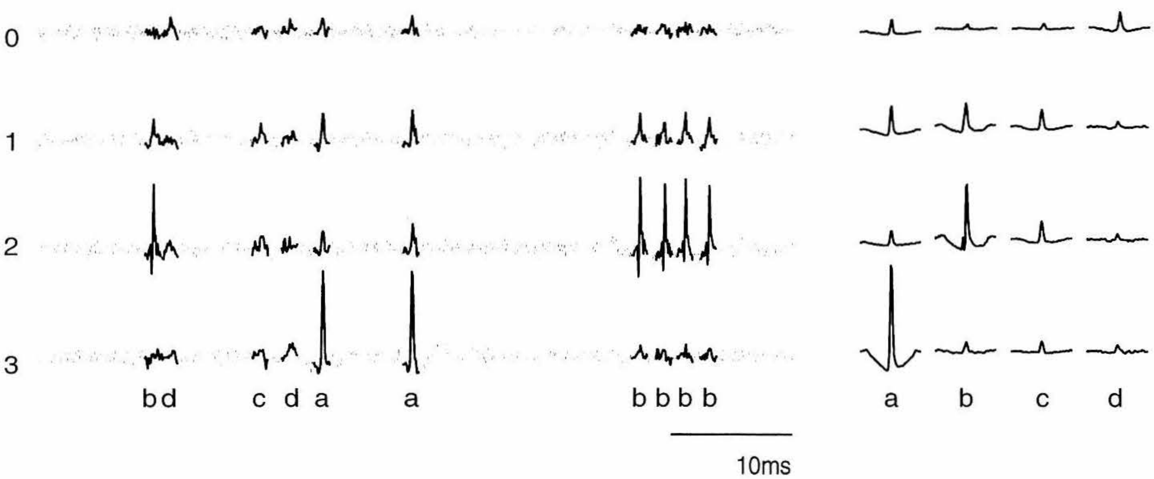

Figure 21: Example Tetrode Recording
On the left is a short stretch of the four channels of data. Heavier sections represent detected events. Letters signify events from four presumably different units. On the right are the averages of all events from each of the identified cells; although not precisely used as such in our spike sorter, these can be considered model spike shapes or templates. Signals have been digitally high-pass filtered at 300 Hz. Horizontal scale is time, and vertical scale is arbitrary scaled volts.

recording only snippets around each superthreshold event, we have the ability to examine the non-spiking portion of the neural signal, the local field, as well as spiking events. The local field has proved a rich source of information, as, being a continuous signal, it can be analyzed even when the cells in an area do not fire to an experimental task, or the stimulus is beyond the receptive field. This is examined, in part, in a forthcoming paper [48].

5.1 Tetrode Signal Spectra

A suite of recordings were selected among those made in two animals (C and L) for good isolation and relevance to a behavioral task not described in this chapter. To broadly characterize the signal, we computed the spectrum of each recording, typically 10–15 minutes in duration, and formed their means and standard deviations at each frequency. The results for this computation are found in Figure 22. The mean spectra from the two animals match well, both showing a broad $1/f$ structure with deviations 5–30 Hz and 200–6000 Hz. In the previous chapter, we examined the upper deviation and showed it to be related to spiking activity. The lower deviation has two potential sources, local field, and mechanical artifact.

Differences in the spectral power between tetrode channels should be related to differences in impedance and differences in signal strength. In example spectra from the previous chapter, spectral changes in the band 300–6000 Hz were informally attributed to changes in spike amplitude. To further examine this, we computed the variance in power across channels for each recording, and

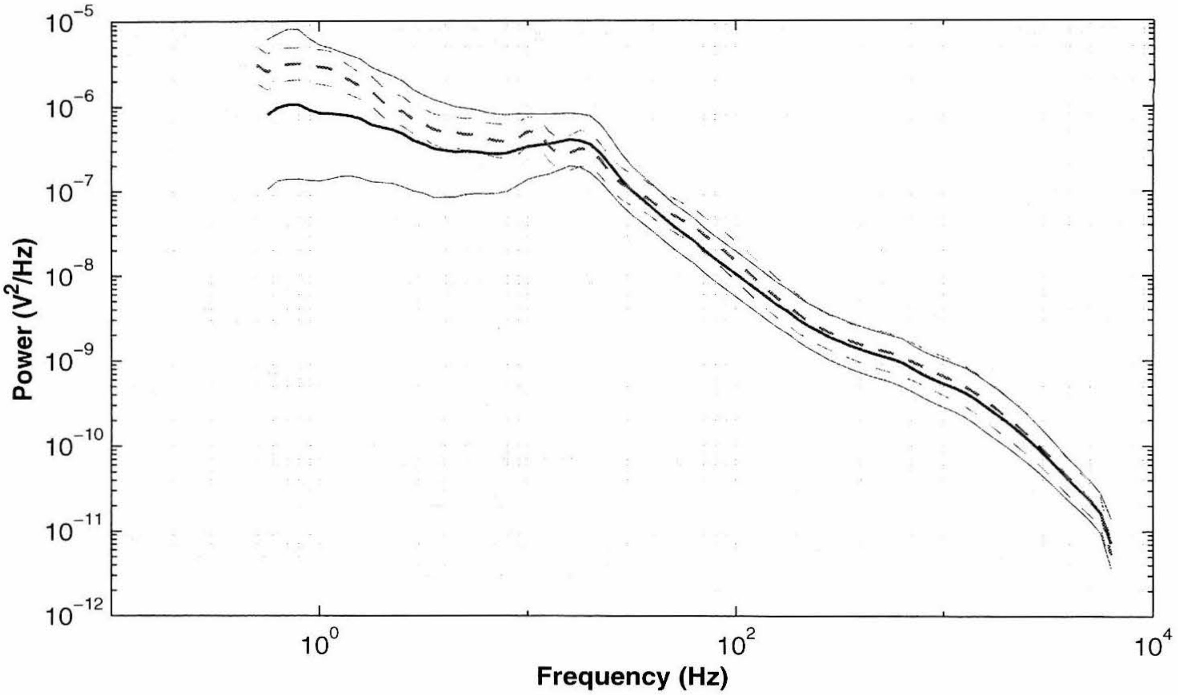

Figure 22: Mean Tetrode Spectra
The mean spectra (heavy lines) from two animals (solid, Monkey C, $n = 43$; broken, Monkey L, $n = 40$) with standard deviations (fine lines). The large variation at low frequencies (below 10 Hz) for Monkey C is due to a change in the characteristics of our recording apparatus when updating from prototype amplifiers to a second generation design, and does not accurately reflect a change in signal.

then averaged the variances. We would expect the spectral variance to be highest for signals with the most local source. Results for Monkeys C and L are shown in Figure 23. The variances are highest for a Gaussian-like peak centered at about 1000 Hz, and not quite completed at the upper end by the recording cutoff. Variance also increases at the lowest frequencies, but, importantly, is low for the frequency band 10–100 Hz where we expect to find local field effects.

5.2 Signal to Noise

The exact definition of signal to noise is difficult for signals that are akin to delta functions. We will therefore report two ratios for signals that have been high-pass filtered at 300 Hz. In both cases, the mean event will be computed for all events that cross a 4 × RMS value, where the RMS value is computed for the entire recording, and events are taken to be 3 ms long, 1 ms before the crossing, and 2 ms after the crossing. The first ratio will be the ratio of the RMS of the mean event to the RMS of the recording, and the second will be the peak value of the mean event to the RMS of the recording. In both cases, each computed data point will be the mean across the four channels for a given recording. The results are shown in Figure 24.

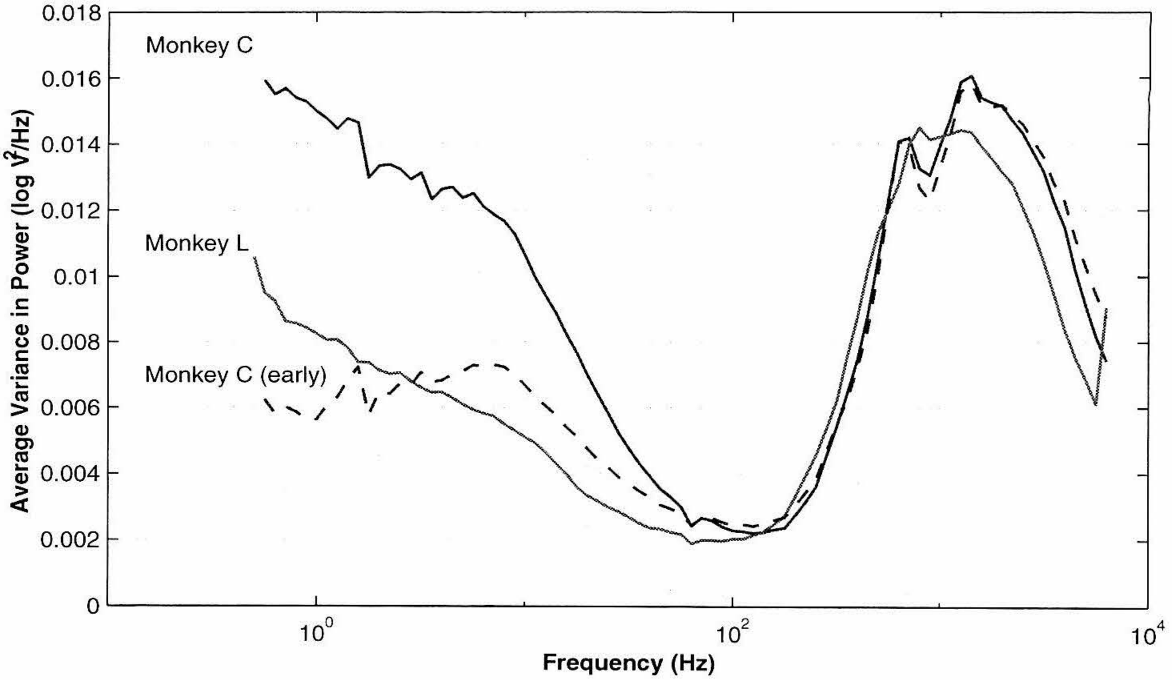

Figure 23: Spectral Power Variance

The mean inter-channel variances from the two sets of recordings. For each recording, the variance in spectral power across channels was computed, and then the sets combined to form the mean. The recordings made for Monkey C (n = 43) contain a change of equipment which affected the lowest frequencies; the average variance before the change in equipment (n = 35) is shown as a dashed line. All of the recordings for Monkey L (n = 40) were made before the change.

The mean RMS-to-RMS value is 1.78 ± 0.46 and the mean peak-to-RMS value is 5.75 ± 1.56. Both figures will be biased upwards by the $4 \times$ RMS detection threshold. The first value indicates that there is almost twice as much power in a spike, over a 3 ms period, as compared to the mean recording. The second value indicates that the mean maximum value attained for each spike is almost six standard deviations out from the noise process; combine that with the expected continuity of spike waveforms, and we conclude that spikes are clearly identifiable and form a process separable from the background noise.

5.3 Penetration Straightness and Tissue Insult

We have histologically identified all marked penetrations and have determined that the tetrode bundles ran straight and the tracks are not unusual. The deepest tetrode penetration, some 12 mm, included 11 lesions which lie in a straight line, even at the maximum extension, as can be seen in Figure 25. Informal examination of tracks from control penetrations made with traditional electrodes and those from the tetrodes shows little difference between the two. That is, the tetrode bundles caused no additional tissue insult.

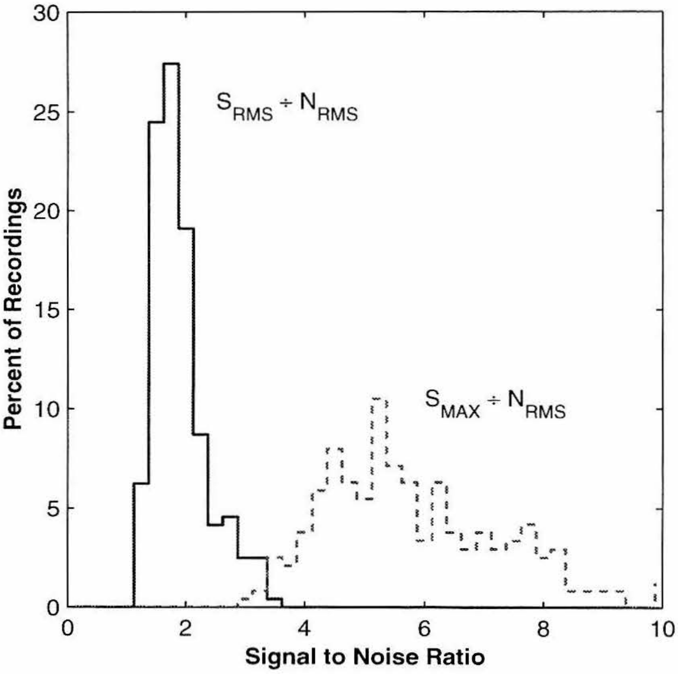

Figure 24: Distribution of Signal-to-Noise Values
Two histogram distributions of the RMS of the mean event normalized by the RMS of the entire recording (solid), and the peak of the mean event normalized by the same quantity (dashed), computed for each recording. Sets of distributions across animals were not significantly different (2-tailed t-test, $p < 0.01$*), and were therefore pooled for this figure.*

An example reconstruction can be seen in Figure 25 showing the straightness of penetrations. This false-color image is an overlaid montage of 11 sections, subjectively aligned to background features rather than the lesions. Nissl stained sections were produced with standard techniques by Dr. Ken Grieve, then photographically scanned, and digitally manipulated in color and position, to generate the figure, by the author. Notice that in this penetration, the tetrode has apparently passed through a small area of vasculature near the second lesion from the top without deviating from a straight path.

6 Discussion

During development of this technique, we tried various alternatives for establishing electrical contact from the wires in a bundle to the individual connector pins on the carrier. Some researchers recommend using a low-temperature alcohol or butane flame to flash the insulation off the tips of the wires. We found this left the wires brittle, and thus prone to breaking, while inspection under a light microscope revealed not nice clean metal, but a wire encrusted with black flaky material. Alternately, other researchers recommend careful scraping of the insulation to clear it away. We found this to be difficult and yielded mixed results for insulation removal, while nearly always curling the ends of the

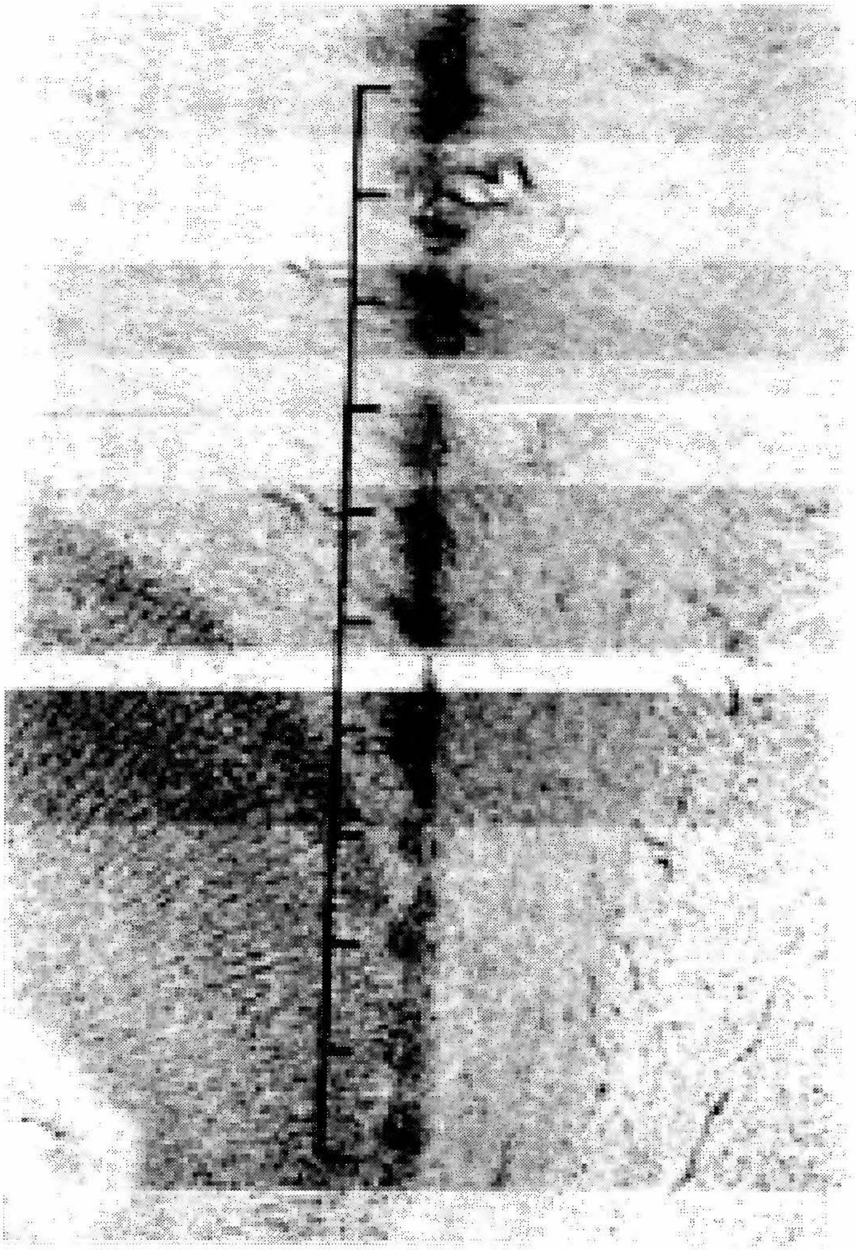

Figure 25: Series of Lesions
11 lesions left along a tetrode penetration. Black ticks are 500 μm apart.

wires, making them frustrating to handle. Our first alternative was to do nothing: connections were made by directly immersing the tetrode wire in conductive paint without any effort to remove the insulation. This, remarkably, worked. While our initial report [54] suggested that the vehicle used in the paint dissolved the insulation, we now feel that this is unlikely, and instead that contact was being made solely through the exposed end of the wire. Our second developmental stage was to use the chemical stripping agent as described above which, while slower, leaves a length of clean fresh exposed wire. After the transition to this technique, the measured tip impedances dropped, and the signal improved, presumably because of the elimination of the high-resistance contact. Given sufficient time stripping time, this technique has proved entirely reliable, and failures due to improper electrical contact are, at most, rare.

7 Summary

We have described the design and construction of a tetrode design appropriate for use in the chronic monkey preparation. Experiments were performed to verify the mechanical and electrical performance of the tetrodes which showed them to be functional and adequate for simultaneously detecting multiple neighboring cells in monkey cortex.

And I follow the tracks that lead me down
— SARA MCLACHLAN (Black, *Solace*, 1991)

Chapter Four: Chronic Chamber Microdrive

1 Introduction

The Chronic Chamber Microdrive (CCMD) was designed to allow pseudo-chronic placement of electrodes, specifically tetrodes, within a standard cranially-mounted recording chamber. An application has been made to the United States Patent Office [57] to protect the design of the CCMD, as described in this chapter.

2 Background

Extracellular electrode insertion has traditionally followed one of two paths, the first using chronic electrodes, which are surgically implanted, the second using acute[6] electrodes, temporarily inserted on a daily basis.

Historically, chronic electrodes are not subsequently repositionable: if they are movable at all, it is only in depth. That is, the x and y positions (along a surface tangent to the brain) are fixed during surgery, and the z position (into the brain) is potentially adjustable. In this configuration experimenters can record from a small group of neurons for a long time. However, since chronic electrode insertion is traditionally a surgical procedure, only this one collection, or a small set of such collections, is accessible in a single animal due to the trauma of surgery.

For acute electrodes, a craniotomy is opened surgically, but is then covered by an implanted resealable cylinder, called the recording chamber. Each day, the sealing plug on this chamber is removed and one or more electrodes are inserted in a fresh location. The electrode position is set using an external device called a microdrive (see Figure 26) that mounts to the recording chamber, and adjustable in all three dimensions, but is typically first fixed in x and y, and then scanned in z. At the end of the recording session, electrodes are removed and the chamber re-sealed. In this scheme, the experimenter has access to an larger total number of cells, but can record from a given group of cells for a period measured in hours at most.

We have developed an intermediate solution which allows the insertion of electrodes for a period of many days, but then their retraction and reinsertion in a new location without surgical intervention. The approach depends on a miniaturized microdrive, the CCMD, which can be inserted into the cylindrical recording chamber while maintaining chamber seal (see the figures below). The

---

[6] The common terminology here is regrettably poor, as *acute* is often taken to imply a terminal procedure. For the balance of this chapter, we take *chronic* to mean long-term, that is, spanning an indefinite period of time which may reach weeks or months, and *acute* to mean temporary, that is, spanning a limited period of time which may reach a few hours. *Chronic* also implies multiple recording sessions, whereas *acute* implies at most one recording session.

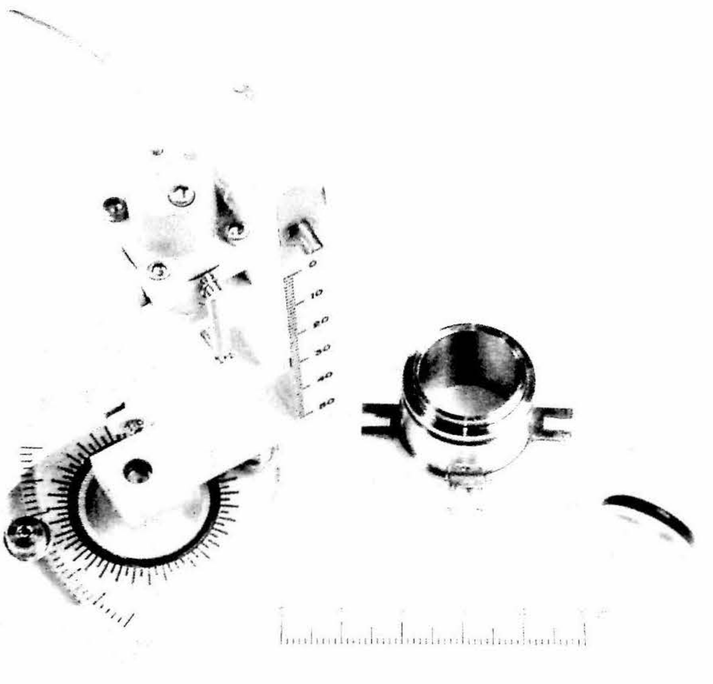

Figure 26: Microdrive, Chamber and Plug
A commercial hydraulic microdrive mounted on a rotational x-y positioner, compatible stainless-steel recording chamber and acrylic plug. See technical report CNS TR-00-02 [50] for a description of the recording chamber and plug. The same chamber appears in later photographs in this chapter.

device holds four electrodes, each individually positionable in depth by a lead screw. The four electrodes are inserted in a group with approximately 200 μm spacing; the exact configuration is not limited by design, but must be chosen for a given microdrive. Using an x-y positioning system comprising two non-concentric cylinders this group of electrodes can be inserted anywhere in the chamber and each one individually advanced in z. After a period of recording, which can span multiple recording sessions, the electrodes can be retracted, a new x-y location selected, and the electrodes re-inserted to depth at a new location, all without surgery. Further, since this microdrive fits within a standard recording chamber, initial exploration can be done with a standard electrode and microdrive and the chronic electrodes inserted only after the target area has been identified, thus improving the yield of the chronic recording.

2.1 More Detail

The CCMD is a plug which replaces the standard, inert, chamber plug, subsuming the normal sealing role it provides, while adding the functions of a microdrive of positioning control and electrical contacts. The combination allows, therefore, for chronic recording in a standard chamber. Seals are provided everywhere by Viton o-rings which show good flexibility, biocompatibility, and imperviousness to standard solvents.

Positional control in the x-y plane is via a dual-non-concentric mechanism. The first degree of freedom is the rotational placement of the CCMD within the chamber, fixed by three externally set screws. The second degree of freedom is the rotational placement of the electrode cylinder within the CCMD, where the two axes of rotation are not colinear, fixed by two vertical clamping screws. This is similar to one of the standard mechanisms used by microdrive manufacturers. Positional control in the z direction is via lead screws, both coarse to adjust the overall vertical position within the chamber, and fine to adjust the vertical position of the electrodes relative to the microdrive. Each of the four electrodes, nominally tetrodes, can be individually positioned in the z direction by independent fine lead screws.

External electrical contact is made through a normal, if somewhat small scale, multi-contact connector, of a type often found in laptop computers, portable telephones, and similar miniaturized electronics. Fine wire electrodes make mechanical and electrical contact with the pins of the connector through conductive paint.

In normal operation, the x-y coordinates are selected and the rotational equivalents determined and set. The CCMD is placed in the chamber and fixed to the appropriate position. Then, the gross vertical control is used to advance the inner parts of the CCMD so that the guide tubes will penetrate the dura. Subsequent to this, the fine vertical controls (the electrode lead screws) are used to advance the electrodes so that they will pass out of the guide tubes and into neural tissue to the desired z coordinate.

The CCMD also has two vertical passages, sealed by set screws, which may be used to flush any void between the bottom of the device and the upper surface of the dura with appropriate liquid, or to apply antibiotics or other treatments. This is especially important for chronic use to combat infection.

2.2 Limitations

The standard chamber has a 19 mm external diameter and 17 mm internal diameter. This is determined by the manufacturer's design (e.g., FHC, Narishige), but is not inherently limiting to our design, except that it constrains the maximum number of individual z positioners. The current design carries four z positioners. This is not an inherent limitation, except against the physical constraints presented by the manufacturer's chamber.

The extent of z positioning is limited by the acceptable height of the device; an arbitrary limit of 5 mm coarse and 15 mm fine was selected.

The extent of x-y positioning within the chamber is unlimited for the reachable range; that is, there is no limit to the resolution for the sufficiently patient experimenter. The reachable portion the the chamber is limited by the need for a dual-concentric design and the chamber o-ring seals. Advances in sealing technology (eg, using a rubberized coating on the sealing surfaces rather than an o-ring) could increase the overall reachable areas. Currently, a ring around the perimeter of the chamber approximately 2 mm in thickness is not reachable; conceivably, this could be shrunk to 1 mm with more advanced materials.

3 Goals

The design problems and goals for this device were as follows. It was first and foremost to be small enough to fit wholly within the standard recording chamber, or very nearly so, with minimal protrusion from the top. The weight, accordingly, was to be limited so as to be reasonable for chronic implantation, while no unusual equipment, such as the balance bars of Legéndy, Salcman and Brennan [34], would be tolerable. It was to carry multiple electrodes, with a lower acceptable limit of two, where each electrode was individually positionable to the greatest extent possible. It was also to provide a means for allowing sufficient access to the chamber for periodic cleaning and application of antibiotics or other treatment. It was to have easy and reliable connections to the electrodes which would not disturb the electrode positions. And finally, it was to be readily protected when not in use to prevent animal sabotage. All of these goals have been met.

4 Solutions

This section will describe the solutions taken to approach the design goals and problems.

The major problem was to allow x-y positioning within the confines of the chamber and without having undue mechanism protruding vertically from the chamber. Although a true x-y stage would be preferable, it was felt the best achievable approach would be to use a dual non-concentric mechanism where one axis of rotation would be the position of the CCMD within the chamber, an the second would be a cylinder within the CCMD, offset from the primary axis. This limits the total available area because of the need for mechanically structural walls and sealing members, but the design manages to make all but an outer annulus of 2 mm available. This is depicted in Figure 27.

Secondary to that was a means to fit multiple independent vertical controls within the chamber. A minimum of two controls was deemed necessary, with four or more being ideal; each control was to positions one electrode. The design manages to squeeze four vertical controls into the inner chamber by using miniature threaded rod and riders. The machining required is not insubstantial.

5 Applications

Normal acute use is possible with the CCMD, as it can be used in the place of a normal microdrive. The x-y position would be selected, the appropriate transformation to rotational coordinates performed, the device adjusted to reflect those positions, and then placed in the chamber and the

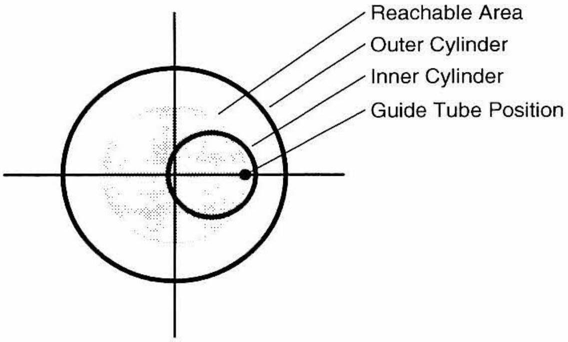

Figure 27: CCMD Reachable Area
The area reachable in the x-y plane using the dual non-concentric positioning mechanism is shown in grey. The outer annulus of the chamber is not reachable using this method because of the need for a sealing member.

electrodes advanced. As the electrode positions are advanced by a hand-turned lead screw, advancement must be slow and sufficient time must be given to allow for tissue settling in the light of stiction problems nominally encountered with smooth electrode motion.

Alternately, a traditional hydraulic microdrive can be used to determine a good set of x-y-z coordinates, and those targeted with the CCMD. This is the normal pattern of use we expect for the device. This method carries the advantages of quick positioning and searching with a traditional microdrive, with the advantages of chronic recording. In practice, the experimenter might take some days or weeks to locate an ideal area for more detailed analysis, and then insert the CCMD, advancing the tetrodes to scan the selected area slowly, with confidence that the electrodes have not moved appreciably from one day to the next.

And, finally, non-invasive methods for localizing the area of interest, such as MRI or ultrasound [20] can be used to guide CCMD placement.

6 Equations

Adjustment of the dual non-concentric CCMD mechanism to target a Cartesian point within the chamber requires appropriate expressions for conversion between the two coordinate systems. These are given below, in analytic form for CCMD designs which might vary slightly in size from the one described, and numeric form for the presented design.

The forward conversion from angular positions of the CCMD to Cartesian equivalents for the guide tube hole achieved as a straightforward series of translations and rotations. Referring to Figure 28, we define terms: radius r_1 is the distance from the outer cylinder center to the inner cylinder center; radius r_2, the distance from the inner cylinder center to the guide tube hole; angles α_1 and α_2, the angular positions of the two cylinders; Cartesian coordinates x and y, the guide tube hole position. To transform from angular to Cartesian values, we start at the origin, translate

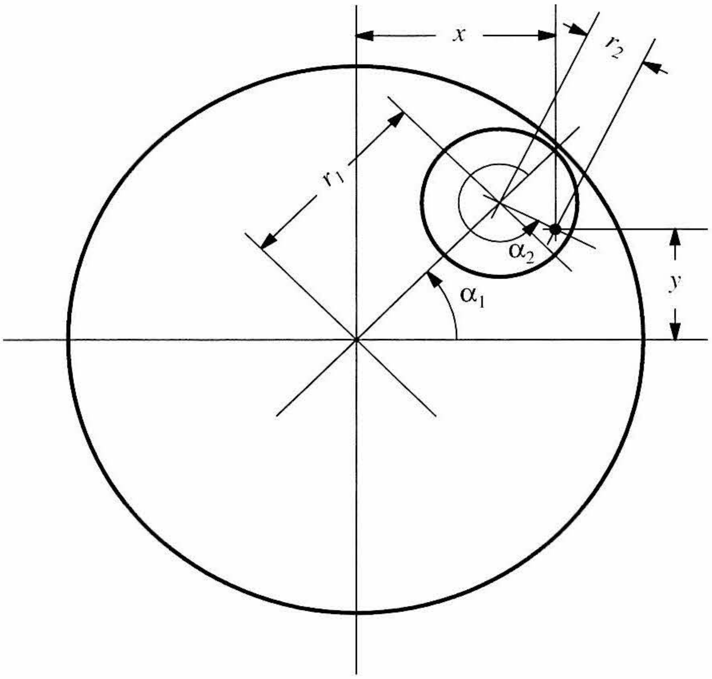

Figure 28: CCMD Rotations
The angles and distances to convert between rotational and Cartesian equivalents. The larger circle represents the outer cylinder, the smaller circle the inner cylinder, and the dot the guide tube hold. This diagram is not to scale, and is misleading for the presented design in that the placement of the guide tube hole does not reach the center of the outer cylinder. Nevertheless, it is a useful representation for the general case of dual non-concentric mechanisms.

to $(r_2,\ 0)$, and rotate about the origin by α_2, corresponding to the guide tube hole's displacement in the inner cylinder, and the inner cylinder's rotational position. Then we translate to $(r_1,\ 0)$, and rotate about the origin by α_1, corresponding to the displacement of the inner cylinder in the outer cylinder, and the outer cylinder's rotational position. The resulting Cartesian coordinates are computed as

$$x = \cos\alpha_1(r_1 + r_2\cos\alpha_2) - r_2\sin\alpha_1\sin\alpha_2 \tag{2}$$

$$y = \sin\alpha_1(r_1 + r_2\cos\alpha_2) + r_2\cos\alpha_1\sin\alpha_2. \tag{3}$$

The inverse problem of transforming from Cartesian to angular coordinates is ill-posed, as there can be zero, one, two, or an infinite number of solutions. Zero solutions occur when the point (x, y) is not reachable with the mechanism; one solution when $\alpha_2 = 0$ for any r_1 and r_2, or $\alpha_2 = \pi$ and $r_1 \neq r_2$; two solutions when $\alpha_2 \notin \{0, \pi\}$; and infinite solutions when $\alpha_2 = \pi$ and $r_1 = r_2$. For conditions with at least one solution, given r_1, r_2, x, and y, we calculate α_1 and α_2 as

$$\alpha_1 = \arctan\left(\frac{y}{x}\right) + k\arccos\left(\frac{r_0^2 + r_1^2 - r_2^2}{2r_0r_1}\right) \tag{4}$$

$$\alpha_2 = \pi + k \arccos\left(\frac{r_1^2 + r_2^2 - r_0^2}{2r_1 r_2}\right), \tag{5}$$

where $r_0 = \sqrt{x^2 + y^2}$ and $k = \pm 1$. The two values of k generate the possibly two solutions corresponding to configurations of the r_0, r_1, r_2 triangle as partially drawn in Figure 28, and as flipped about a line from the origin to (x, y).

In the current design of the CCMD, $r_1 = 0.110$ inches (2.79 mm), and $r_2 = r_1$ (when $r_1 \leq r_2$, the reachable area includes the center of the chamber). With these values, the equations above simplify slightly to

$$\alpha_1 = \arctan\left(\frac{y}{x}\right) + k \arccos\left(0.179 \times \sqrt{x^2 + y^2}\right) \tag{6}$$

$$\alpha_2 = \pi + k \arccos(1 - 0.0640 \times (x^2 + y^2)), \tag{7}$$

where x and y are now specified in millimeters.

7 Sample Recording

Initial *in vivo* testing of the CCMD was done in the first months of 2000. Testing included verifying that when loaded with tetrodes, there were no problems with kinking as electrodes were advanced into clear faux brain (household gelatin) before installation in an animal preparation. When implanted, coordinates were selected according to previously made tetrode recordings in the same animal, guiding the CCMD location to visually-responsive areas. Figure 29 holds a sample CCMD recording from this location, showing all sixteen channels from the four loaded tetrodes.

8 Discussion

The CCMD is neither the only nor the first chronic microdrive. We should, therefore, compare it to to other similar systems, of which there are currently six distinct types in common use.

The hyperdrive first mention by Wilson, McNaughton and Stengel [74], with its various derivations, is commonly in use to chronically position multiple tetrodes in the mouse rat, and cat preparations. Implantation is a surgical procedure and while not moveable in the x-y plane, each electrode is individually adjustable in the z direction.

The Frederic Haer Corporation makes a multiple microdrive intended for acute use. It consists of a collection of eight of their standard microdrives, modified so that they fit closely together and so that the electrodes can all be inserted into the same area. While neither intended for chronic use, nor readily adapted to it, it uses the same kind of dual non-concentric mechanism for x-y positioning, and has independently controllable electrodes.

Another popular acute multi-electrode drive is made by Thomas Recording, based on Reitboek's work [62]. While even the smallest versions of this mechanism are large enough to preclude

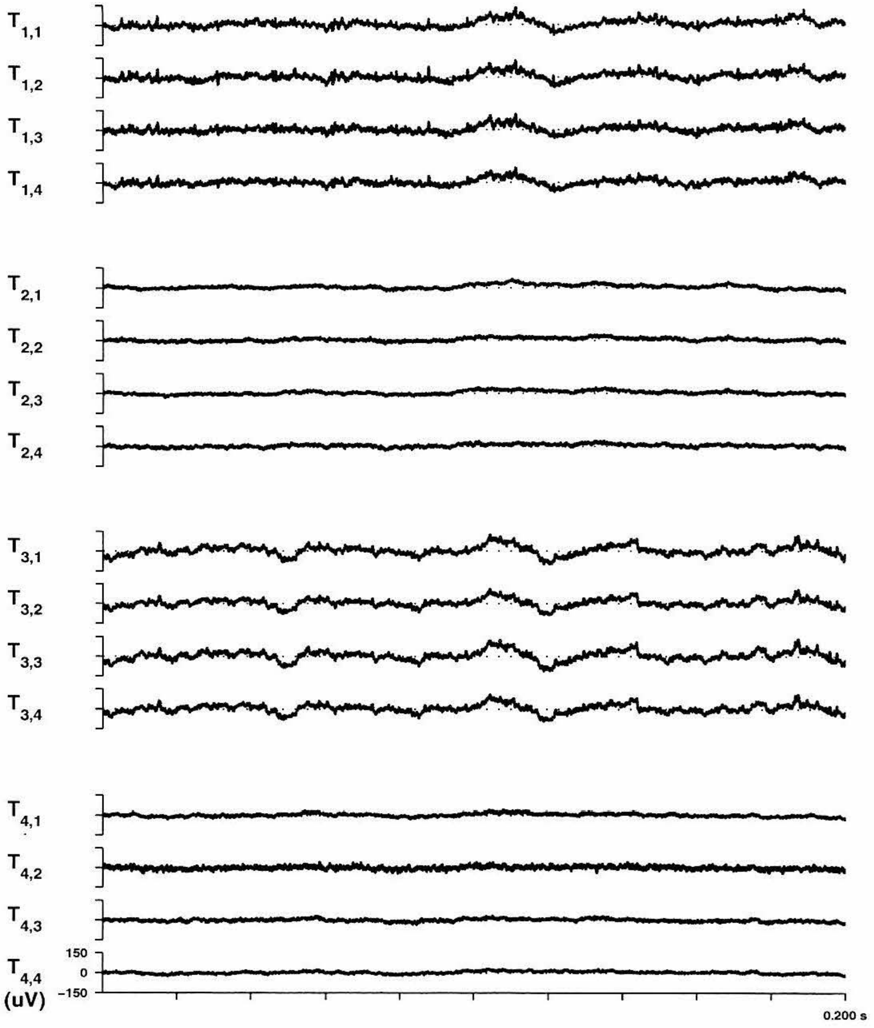

Figure 29: CCMD Sample Recording
Tetrodes 1 and 3 show evidence of spiking and local field activity, indicating they were in gray matter, while Tetrodes 2 and 4 were more likely in white matter.

it's use in the chronic preparation, the drive is stable and easy to use. The electrodes are metal (typically tungsten) wires insulated with quartz that are drawn in a special puller to a thin taper, and are available in monode (1 conductor), tetrode (4), or heptode (7) configuration. The literature contains effusive praise [44] for the qualities of the electrodes, but as yet, a chronic version of the microdrive has not been produced.

Chris de Charms' work [14] might be the closest to such a device. His chronic design uses ultrafine iridium electrodes running in a 350 μm center-to-center grid. Each electrode is individually positionable in the z direction, but moving an electrode from one x-y position to another requires deplantation and reimplantation. This disadvantage is circumvented by populating the entire grid with electrodes in typical use. It would be difficult to adapt this drive to tetrode use unless the quartz fiber electrodes from Thomas recordings were used.

Recent work by Nichols and colleagues [45] described a lead screw based microdrive suitable for use with a Crist grid, and small enough such that three can be simultaneously used in a standard recording chamber. The capabilities of this mechanism are closest to those of the CCMD, while the design is quite different. The screwdrives are fixed to the grid in arbitrary grid position, and the grid may be rotated within the recording chamber. Thus x-y positioning is possible after the electrodes have been retracted, although with a little more difficulty than with the CCMD. Each electrode is individually positionable, and the system can be made to seal the chamber as well. It has the advantage of light weight over the CCMD, while suffering from slightly higher mechanical flexibility.

Michale Fee, with his colleagues, has designed a very small lead screw-based microdrive [71] for the rat, and now finch preparations. This remarkable bit of engineering is quite small. The rat version uses vacuum to pull taut the dura and allow the electrodes to penetrate without guide tubes. The x-y position is fixed at construction, but during implantation can be rotationally positioned providing one degree of freedom. Both versions have 3 mm of independent electrode travel although in principle this could be extended.

9 Conclusions

We have presented the CCMD, a device designed to support a methodological hybrid between traditional chronic surgically implanted electrodes and traditional acutely implanted ones. The device is small enough to fit almost wholly within a standard recording chamber, and controls, in the present design, a ganged x-y positioning of four electrodes, and independent z positioning of each.

10 Drawings

The following pages contain either fold-out or photoreduced mechanical drawings of the design, depending on the publication version. These were used to write the patent application and to create the first copies of the CCMD. Please bear in mind that the electrodes, guide tubes, and electrical connectors are not depicted in these figures.

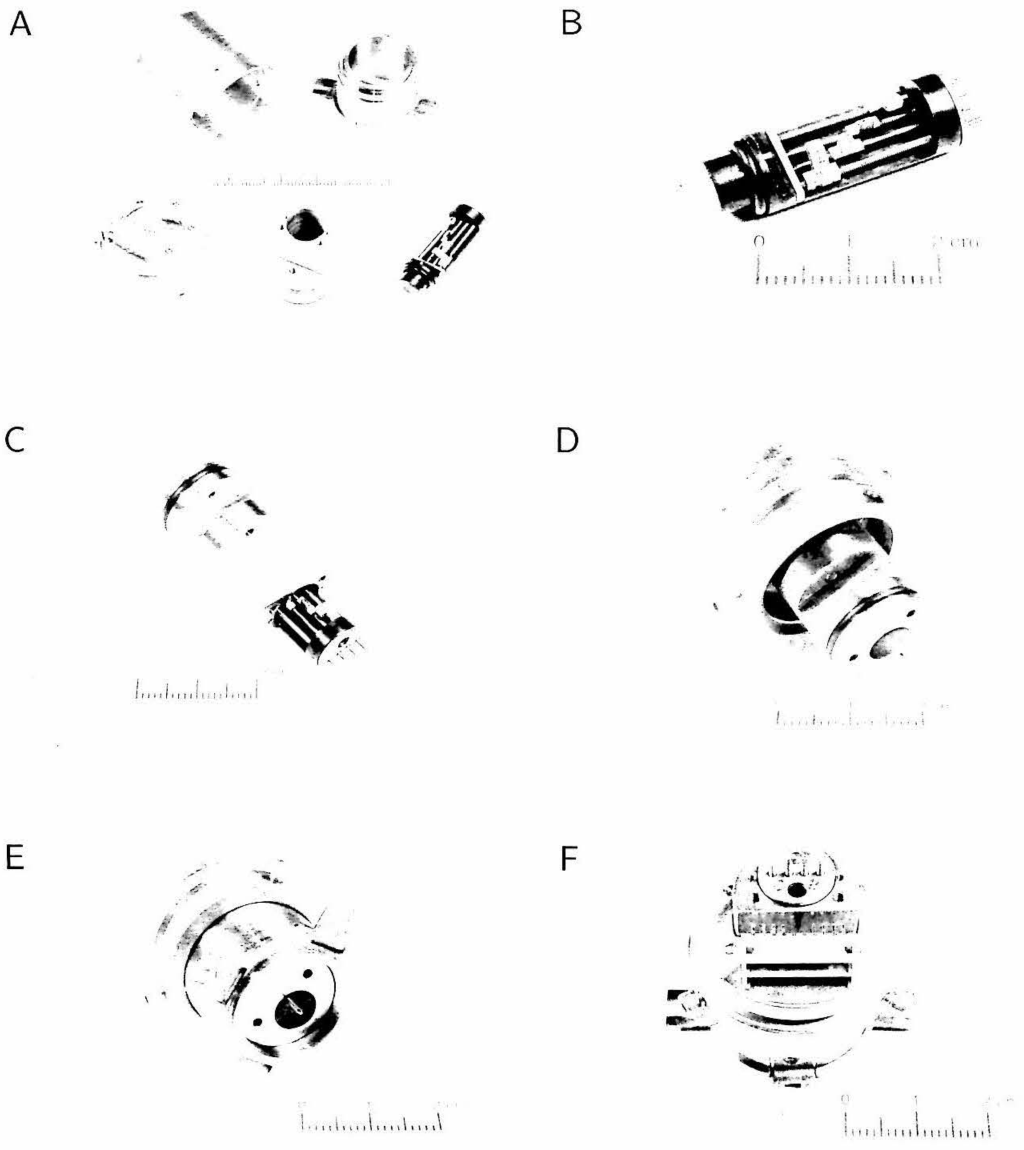

Figure 30: CCMD Assembly
Various views during assembly of the headstage. **A***, the set of parts for a CCMD, disassembled as if just loaded with new electrodes. Clockwise from upper left are the cap, a recording chamber, the inner cylinder, the outer cylinder, and the upper and lower plates (with electrical connectors).* **B***, the inner cylinder with fine lead screws, riders, and guide tube (but without electrodes).* **C***, the inner cylinder being placed into the outer cylinder, with the top towards the viewer.* **D***, the the inner and outer cylinder with upper and lower plates in place.* **E***, the CCMD placed in a chamber, viewed from the bottom.* **F***, the CCMD placed in a chamber, viewed from the top.*

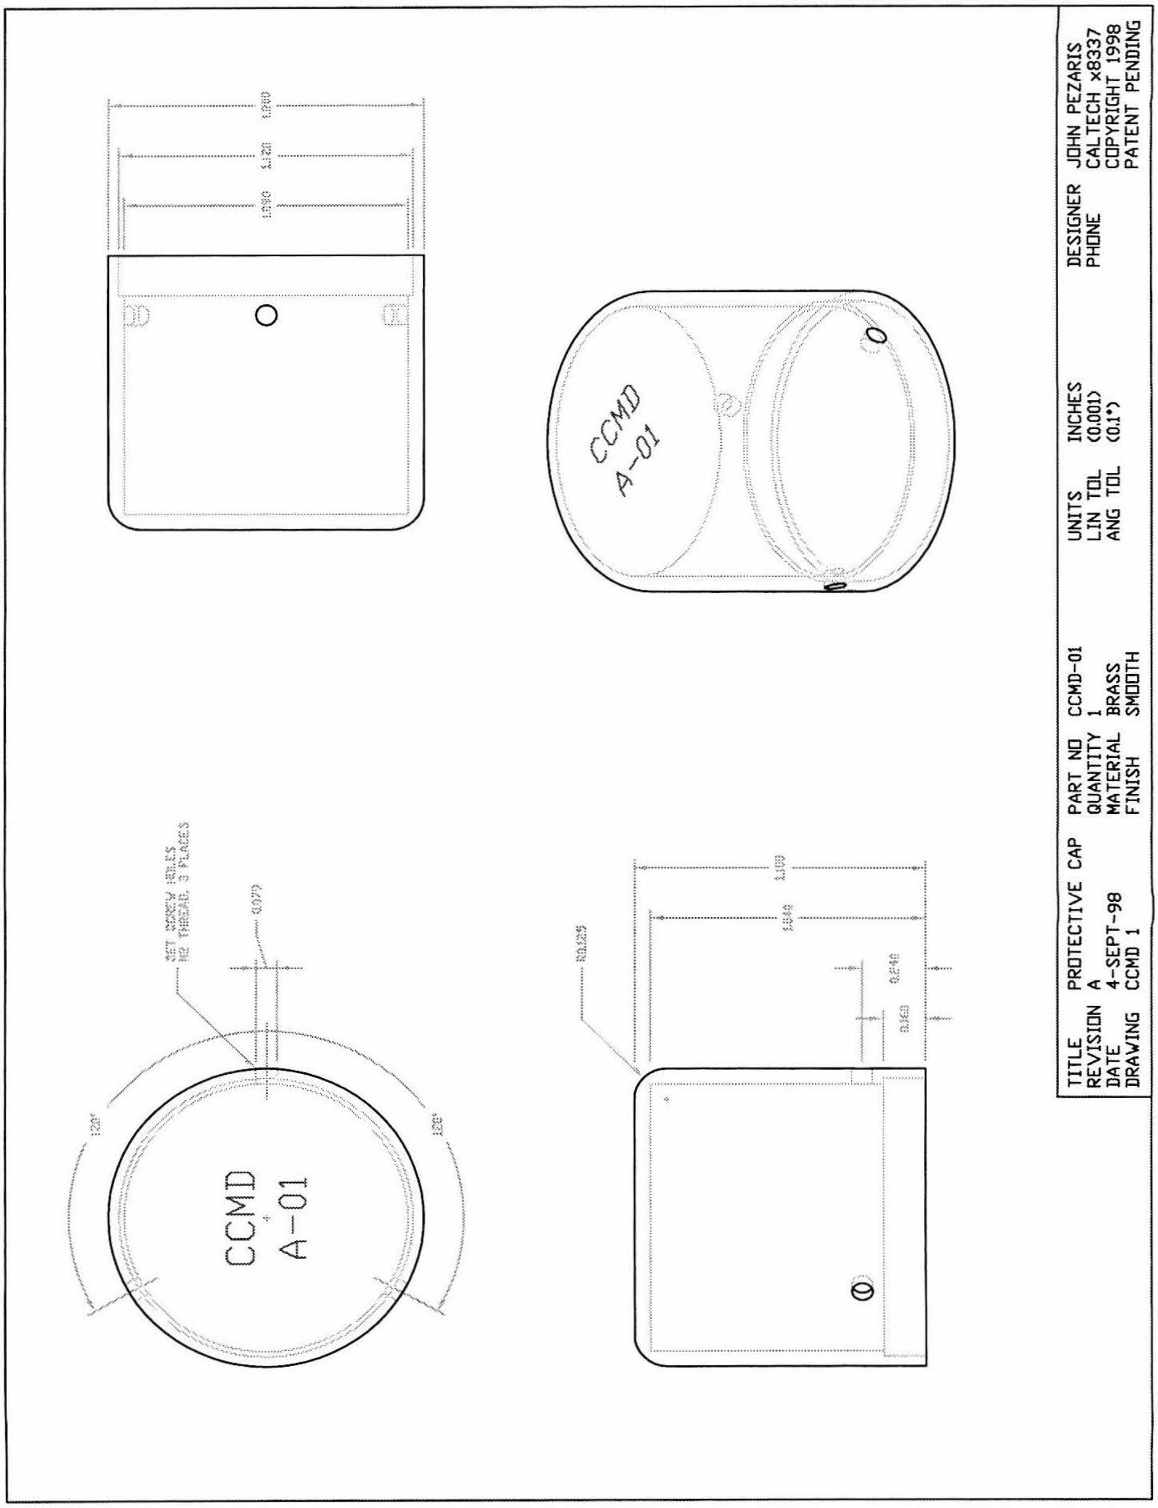

Figure 31: Protective Cap

This cap is placed over the CCMD when it is not in use to protect it from accidental damage or interference.

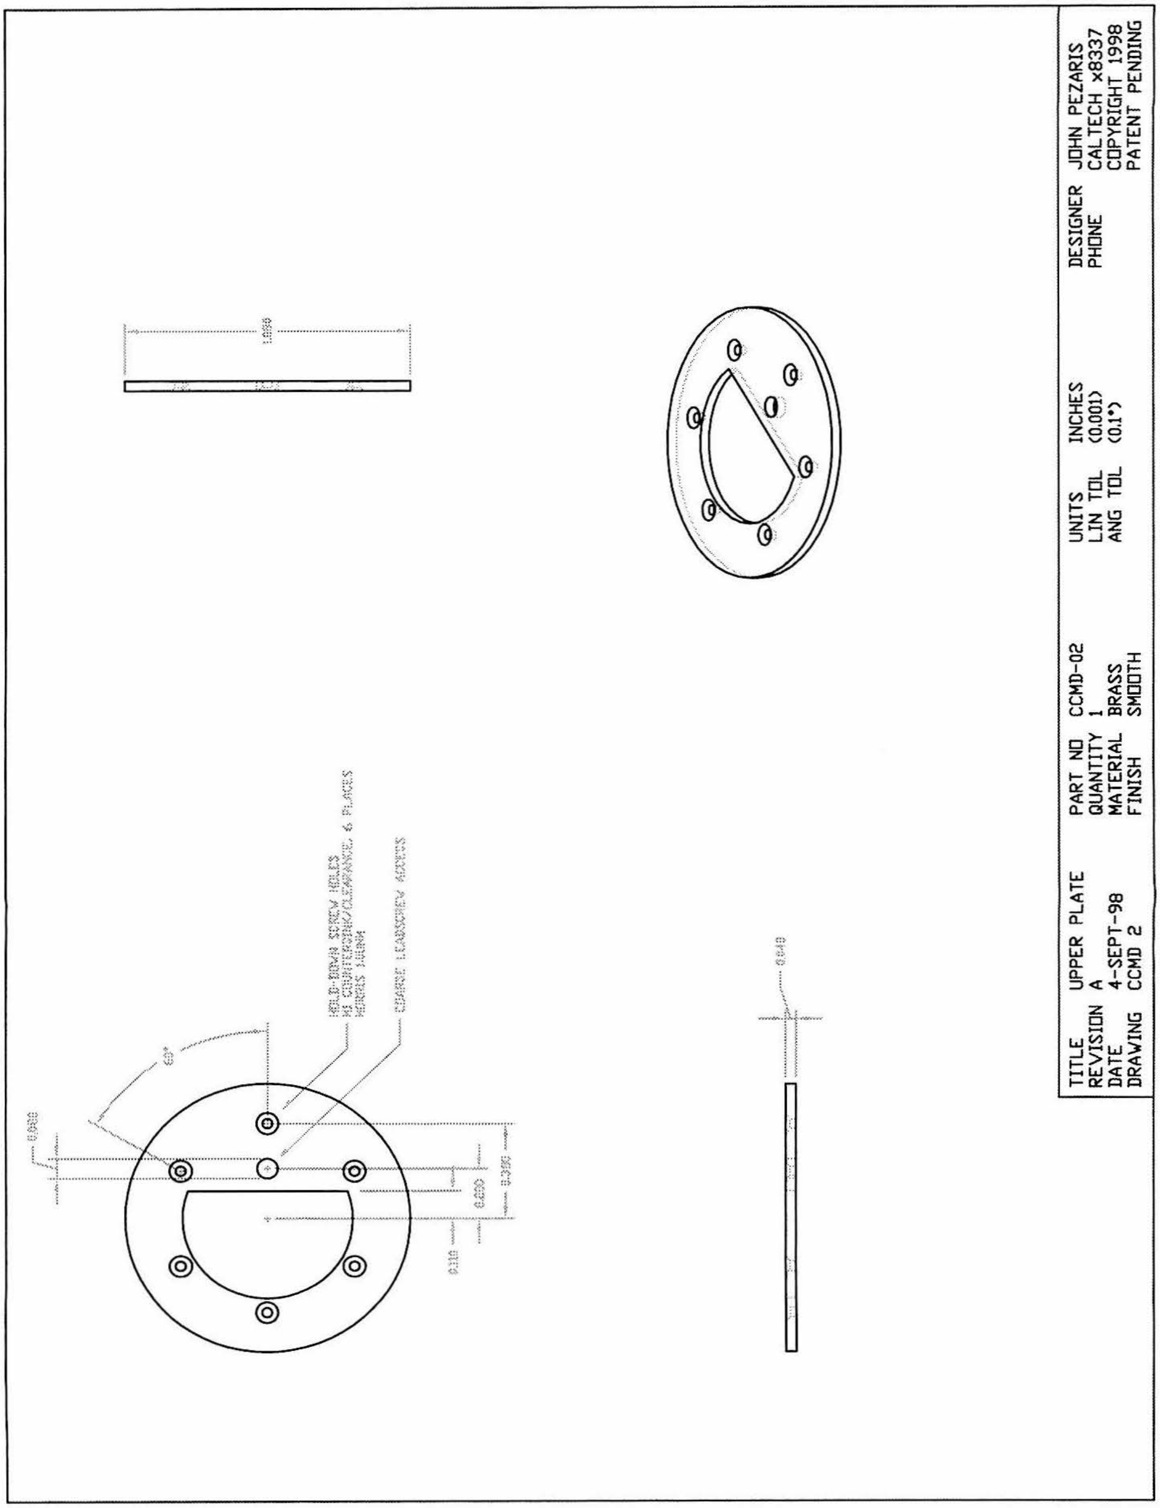

Figure 32: Upper Plate

This plate, along with the lower plate (see Figure 33) forms a well to capture the coarse lead screw.

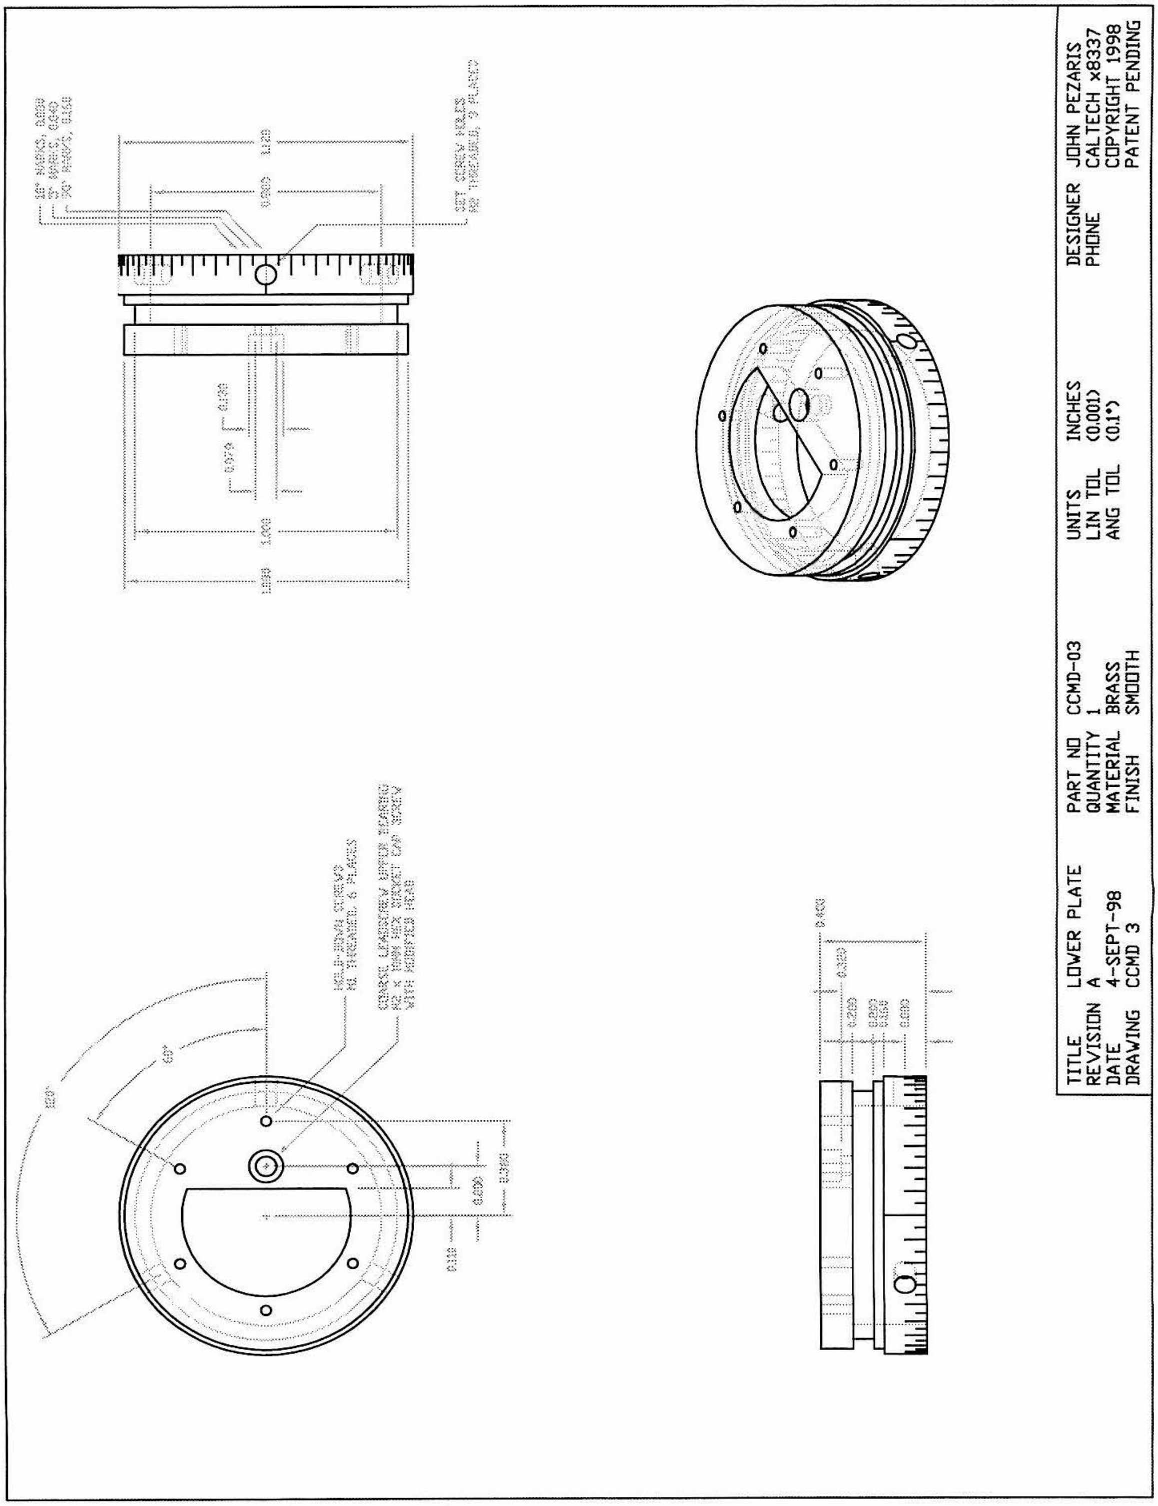

Figure 33: Lower Plate

This lower plate allows the CCMD to be fixed to the chamber. The upper plate is affixed via six screws to the top of this plate, after the coarse lead screw has been placed in the well reserved for it. The outer cylinder rides through the D-shaped hole.

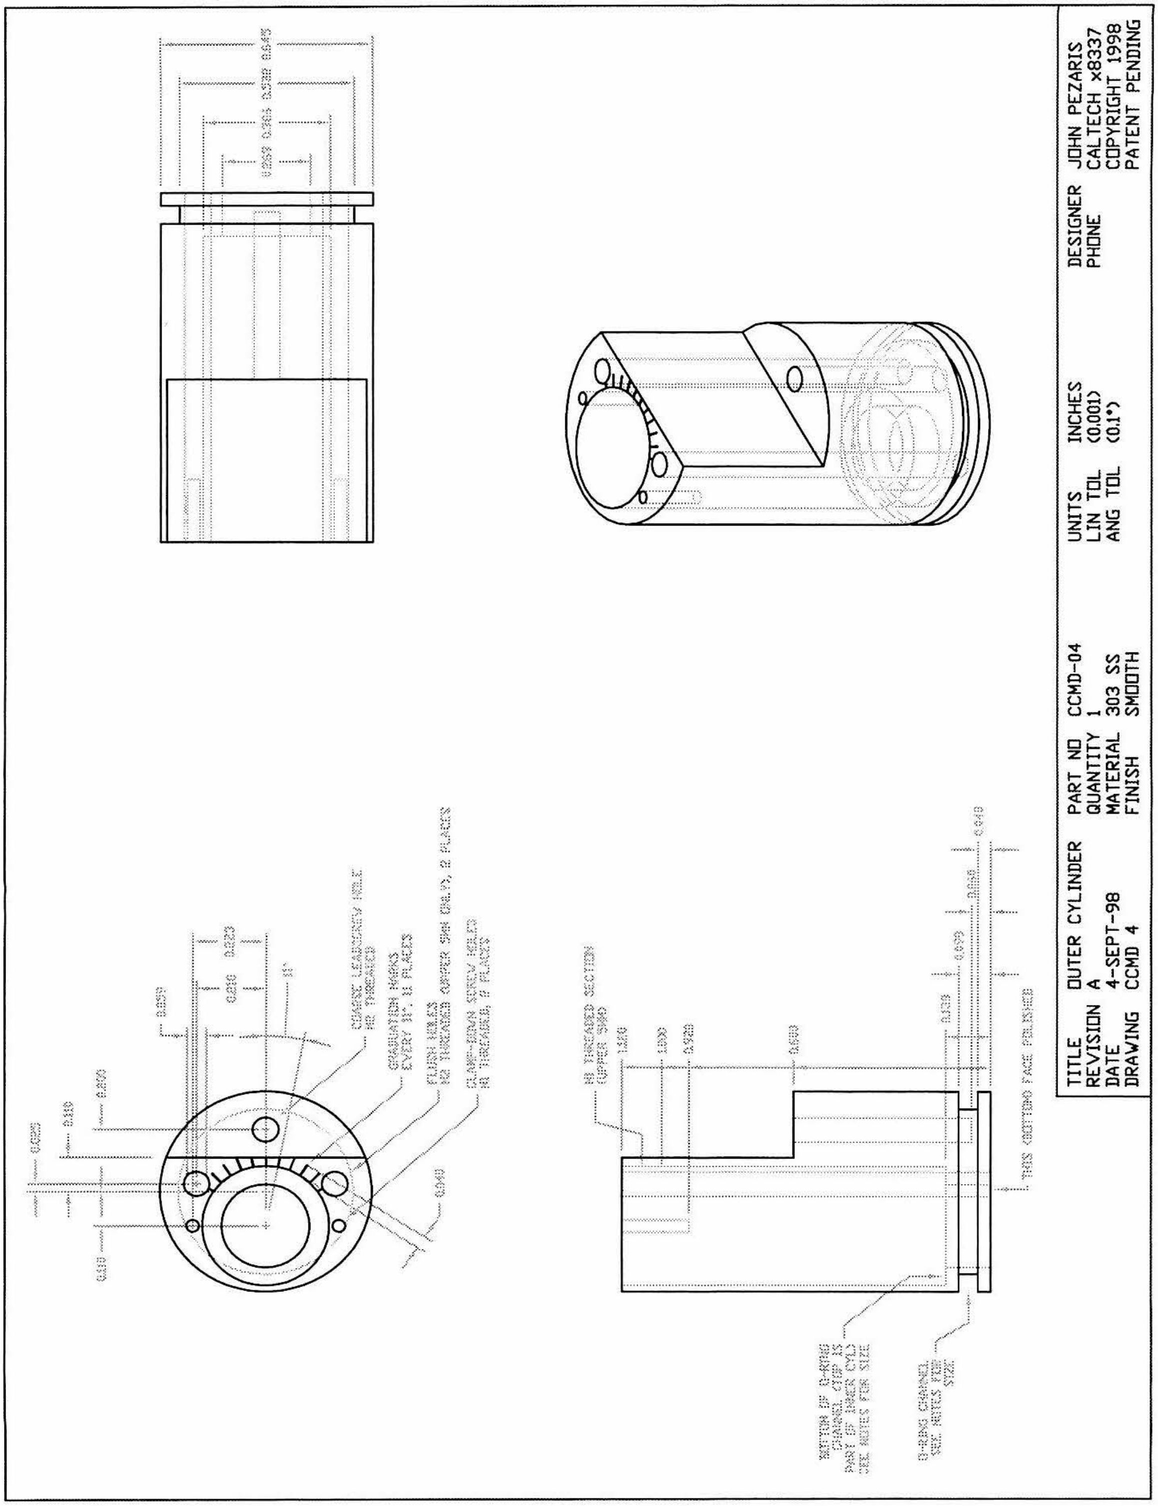

Figure 34: Outer Cylinder

The outer cylinder is advanced via the coarse lead screw, and holds the inner cylinder. There are two holes, threaded at the top, which pass through the outer cylinder to allow access for fluid and antibiotics to the dural surface.

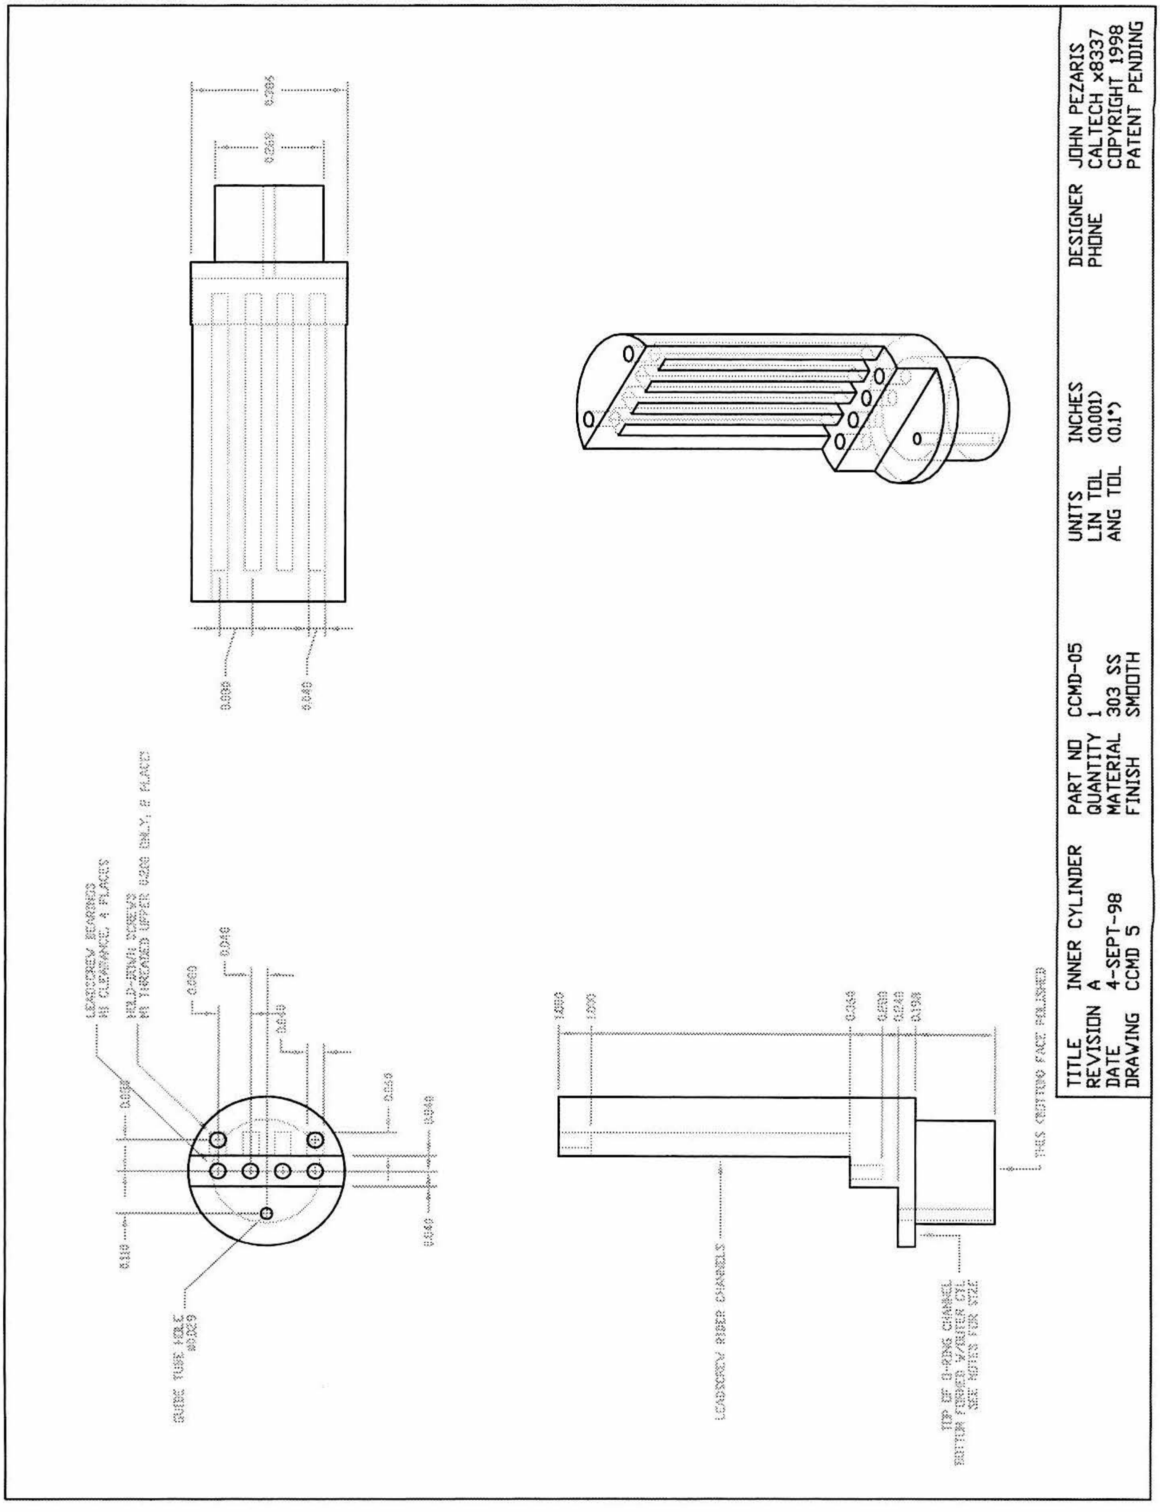

Figure 35: Inner Cylinder

The inner cylinder holds the guide tube in the small offset hole at the base, the four fine lead screws, and the four lead screw riders. The riders run along the four vertical channels which prevent them from rotating, and enforce vertical travel for rotation of the lead screws.

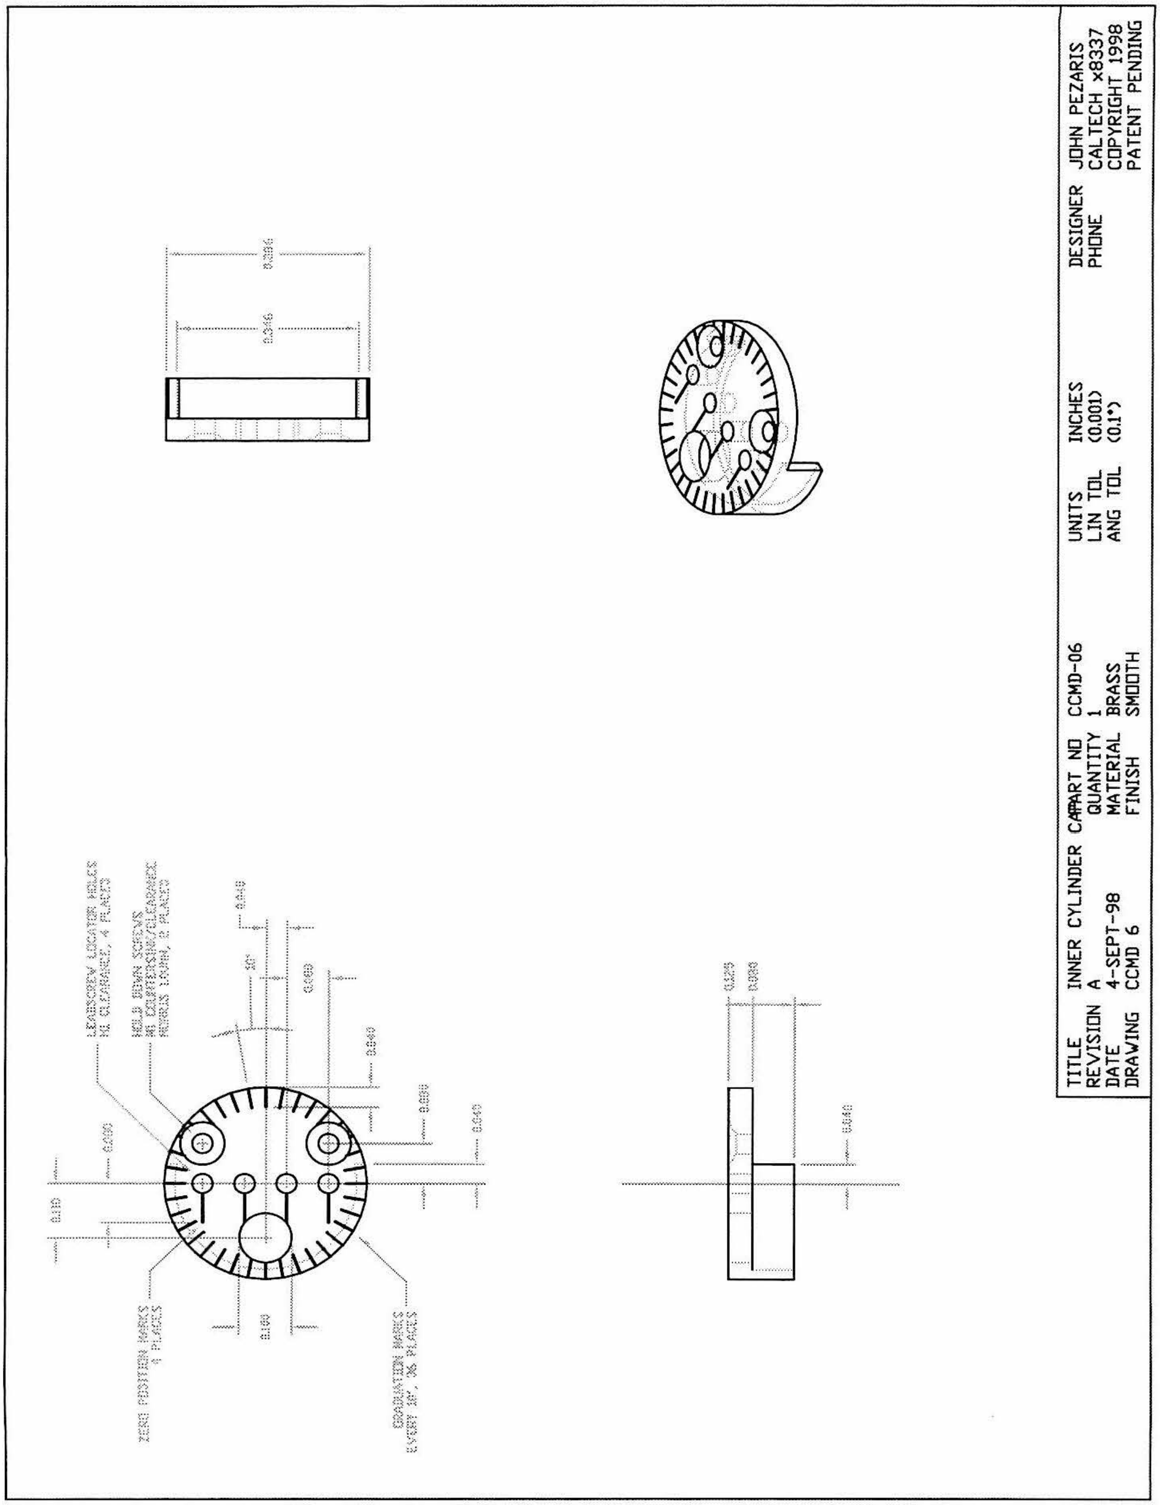

Figure 36: Inner Cylinder Cap

The inner cylinder cap mounts on the inner cylinder, forming a thrust bearing which captures the four fine lead screws. The large offset hole is for electrodes to exit the inner cylinder, on their way to the electrical connector.

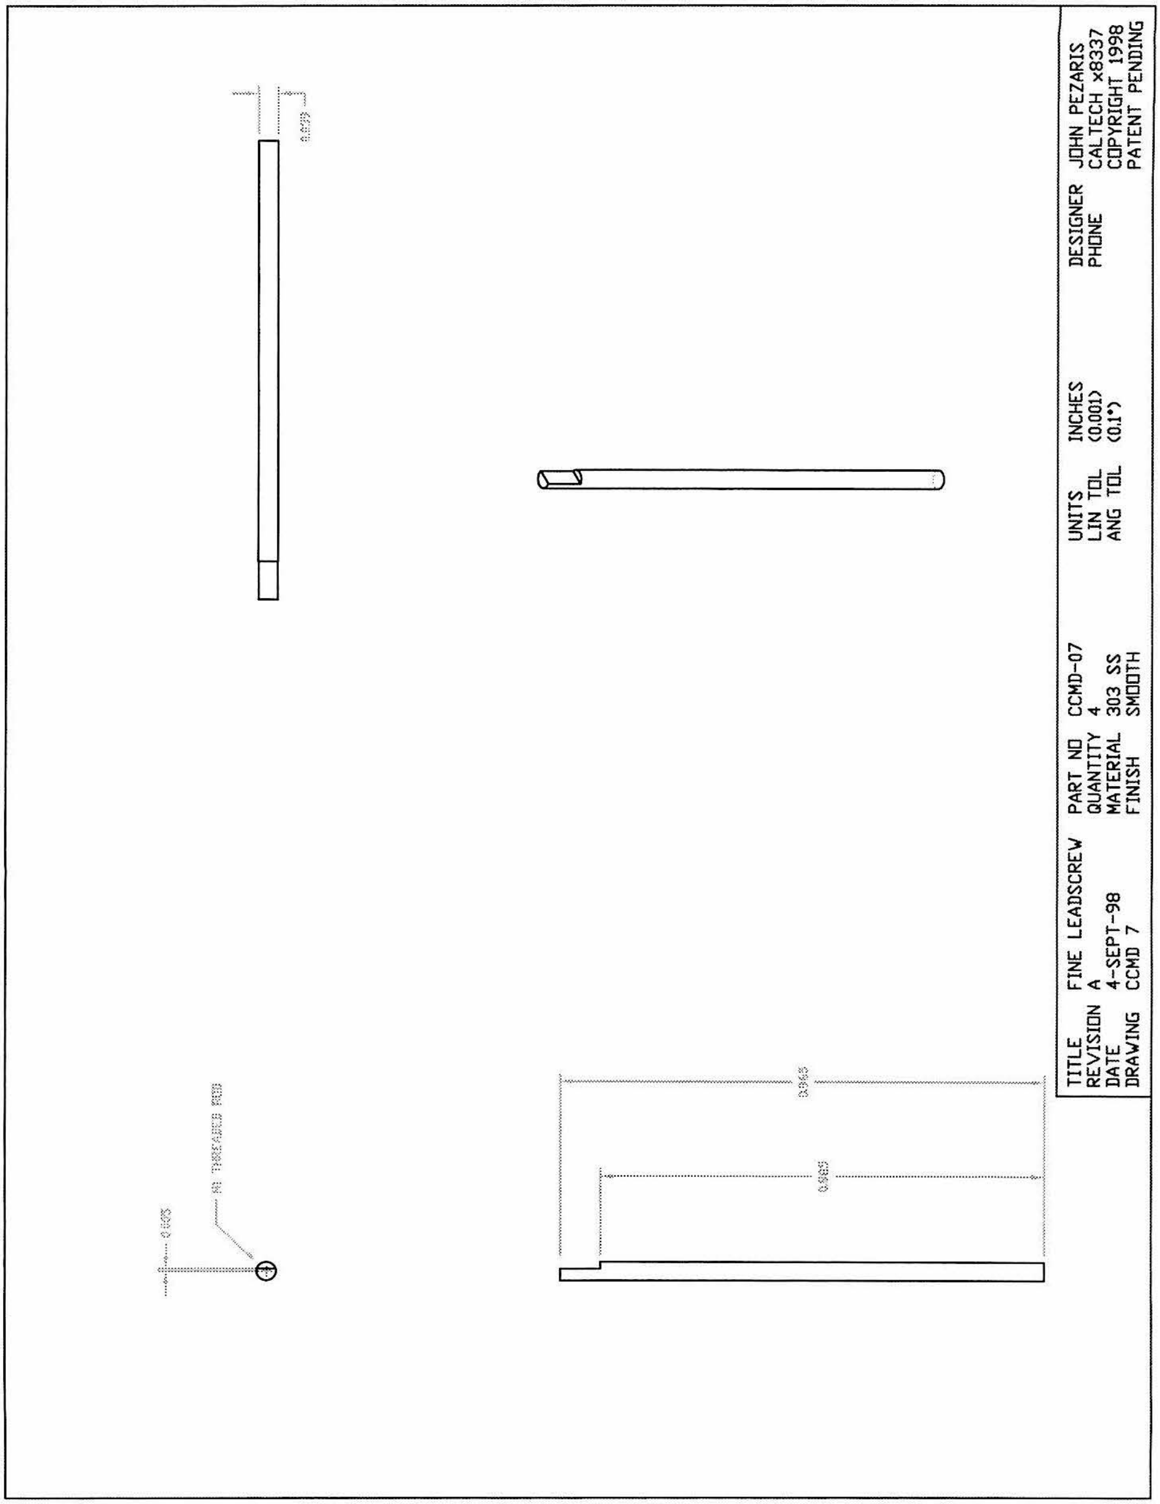

Figure 37: Fine Lead Screw
The fine lead screw is a 1mm diameter threaded rod, shaped at the top to allow rotation with a tool. Four such lead screws fit into the inner cylinder, each having a separate rider.

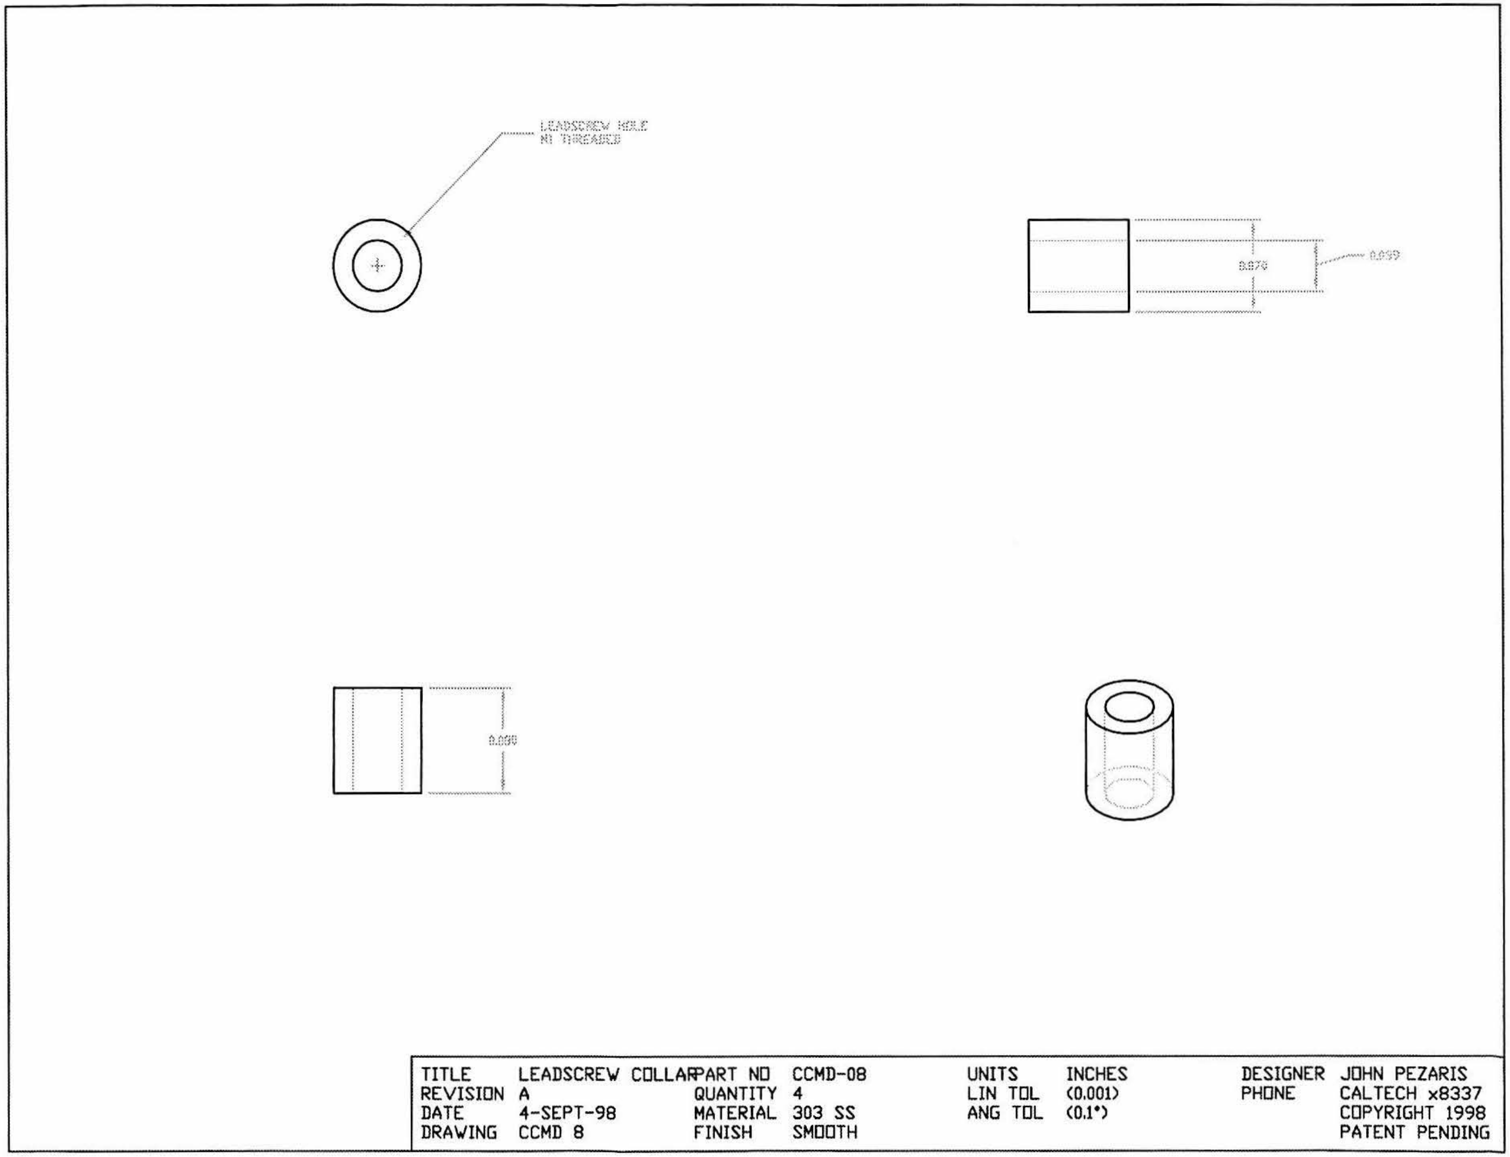

Figure 38: Lead Screw Collar

The collars are glued to the lead screws using a thread-locking compound such as made by Locktite, and form the second halves of thrust bearings against the inner cylinder cap, vertically locating the fine lead screws.

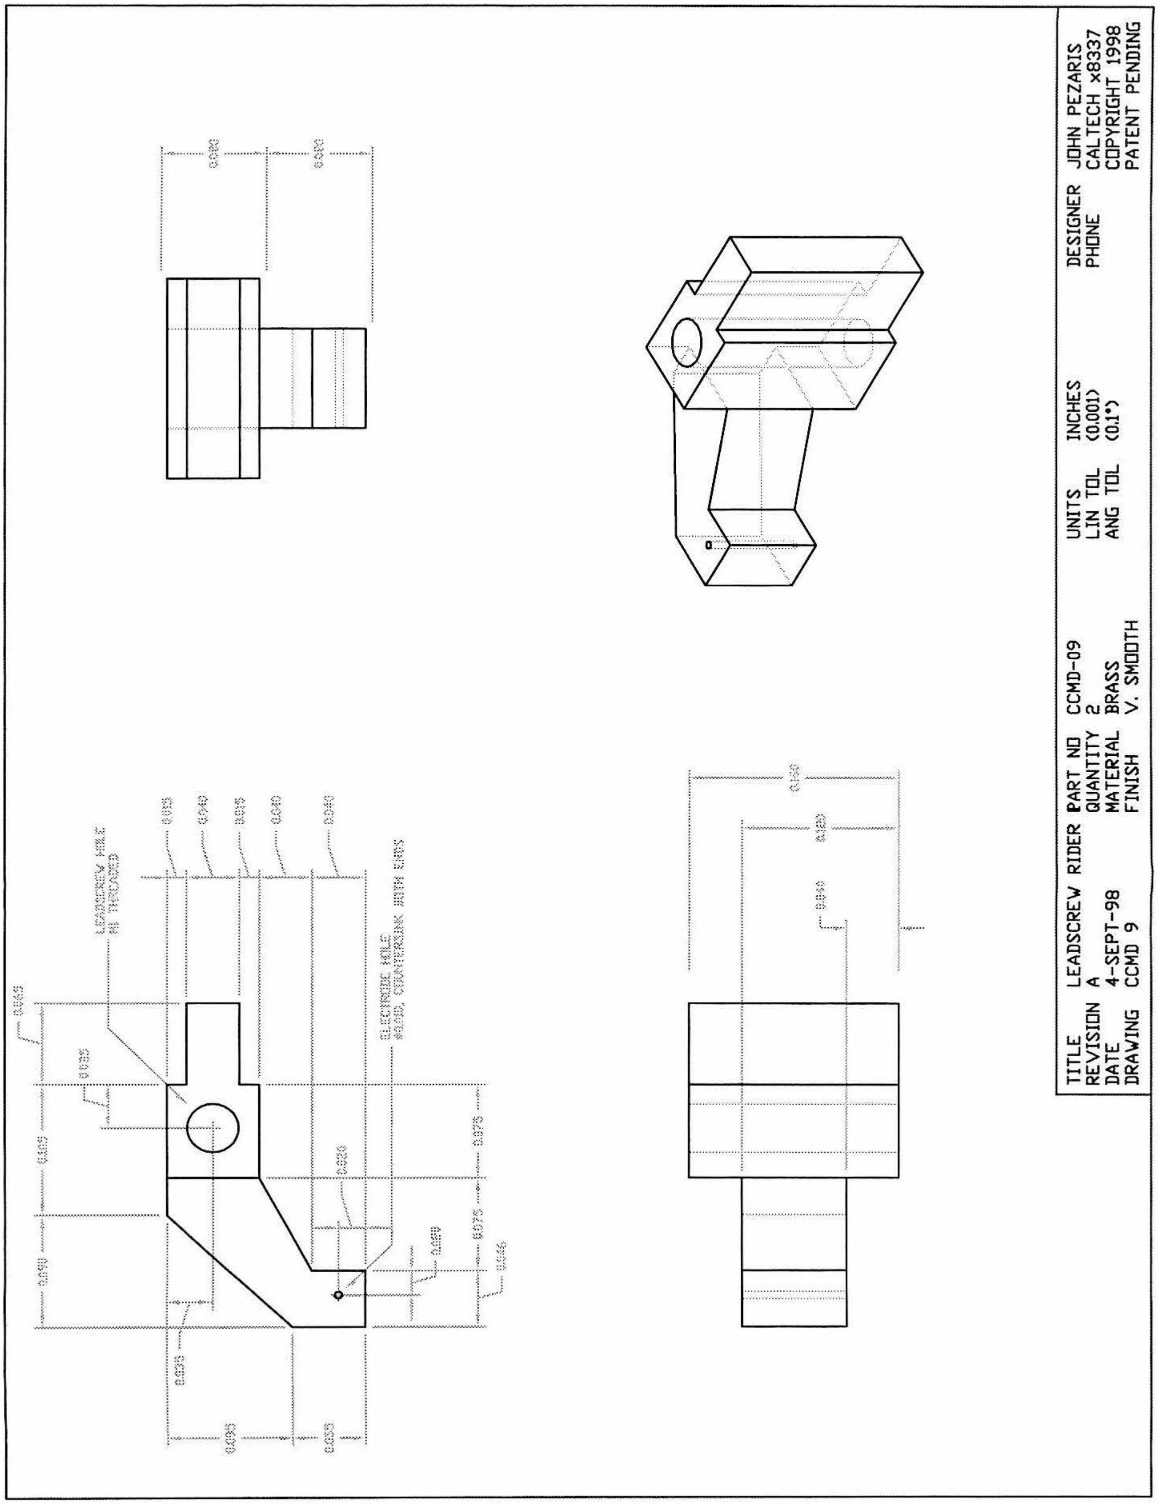

Figure 39: Lead Screw Rider 1

Both lead screw riders are designed to be used in two positions, one placed upside down relative to the other on the lead screw. The outer two lead screws being this larger design. The fine hole is used to hold an electrode, affixed with a drop of removable glue.

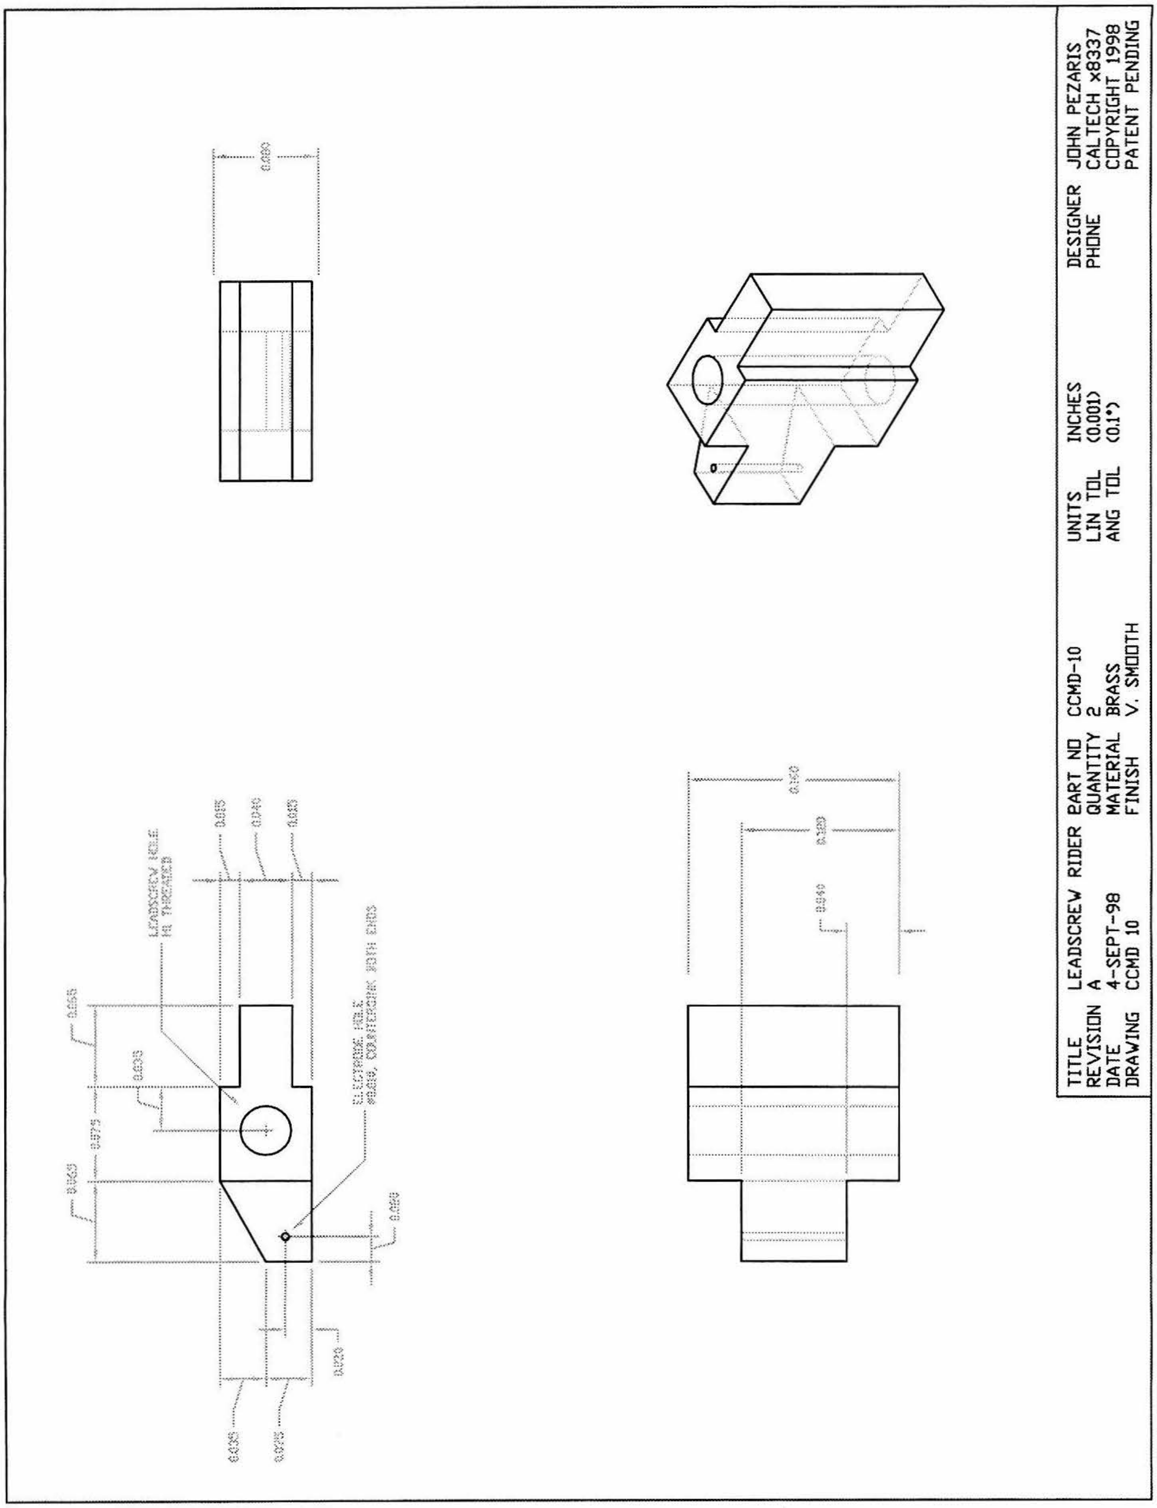

Figure 40: Lead Screw Rider 2

The inner lead screw rider design. This one is, in both positions, used for the two inner riders. The two designs, rider 1 and 2 are made to have sufficient clearance to allow independent unimpeded motion along the entire lead screw.

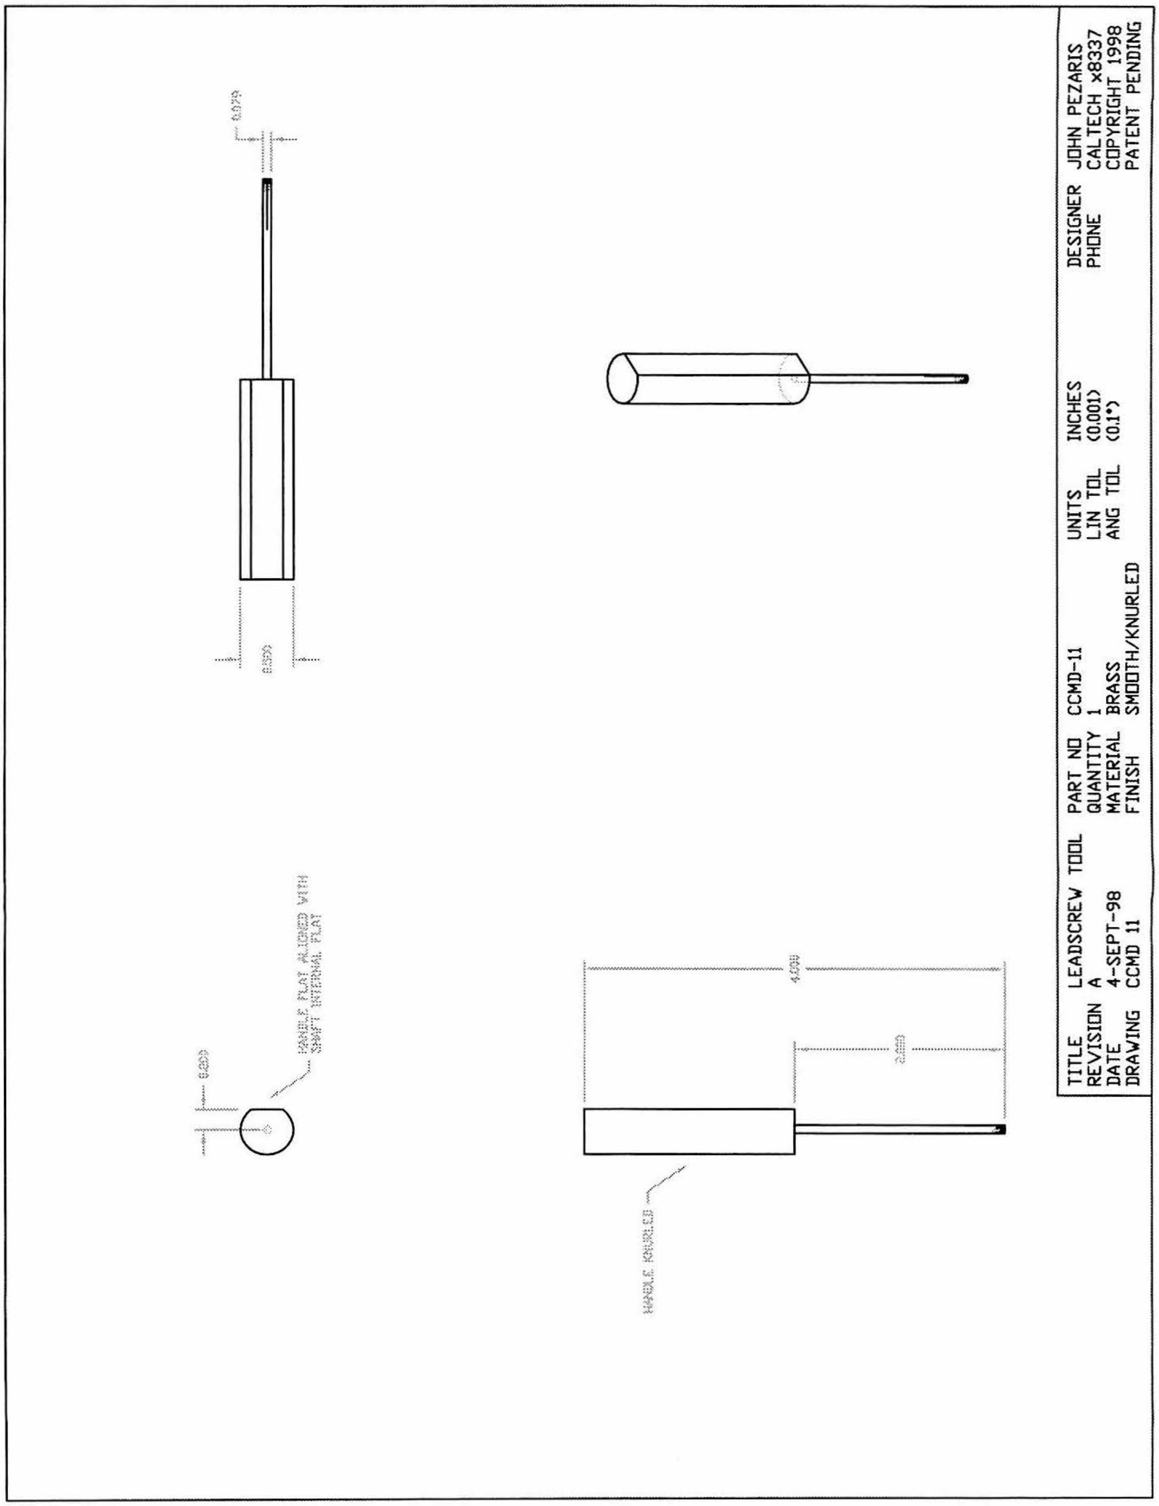

Figure 41: Lead Screw Tool
A tool to advance the lead screw in controlled fashion. Each rotation of the lead screw corresponds to 250μm of vertical motion.

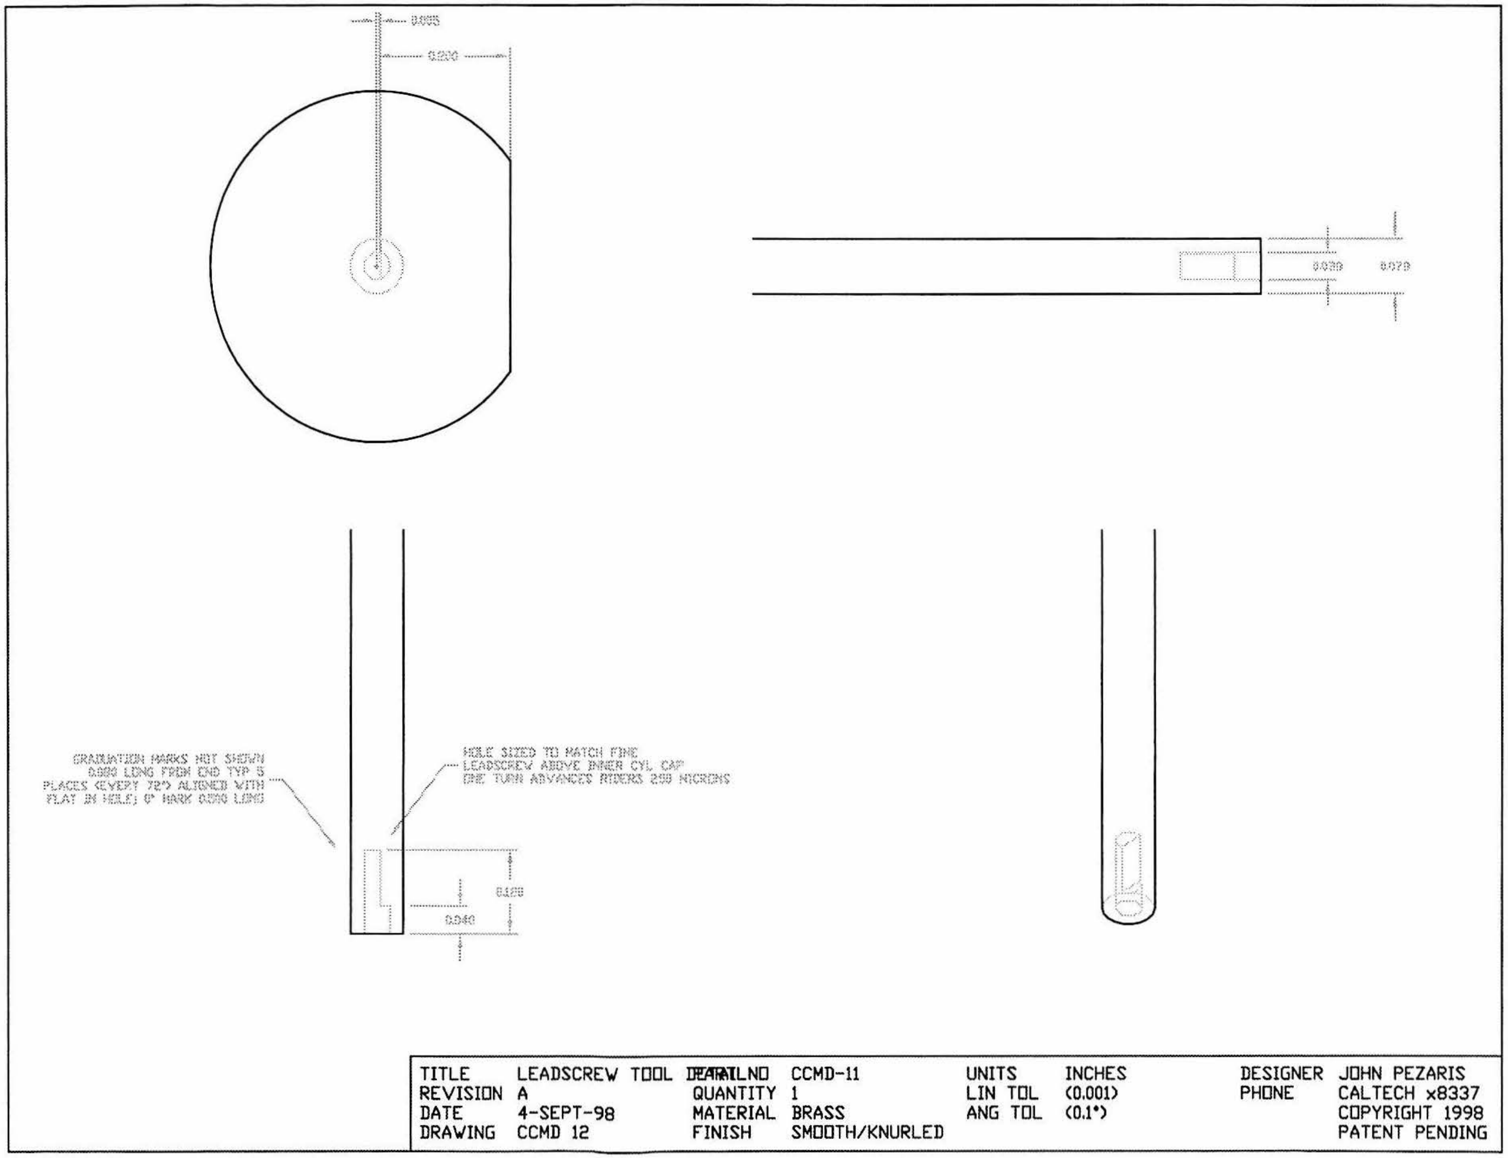

Figure 42: Lead Screw Tool Detail
Detail of the operational end of the lead screw tool.

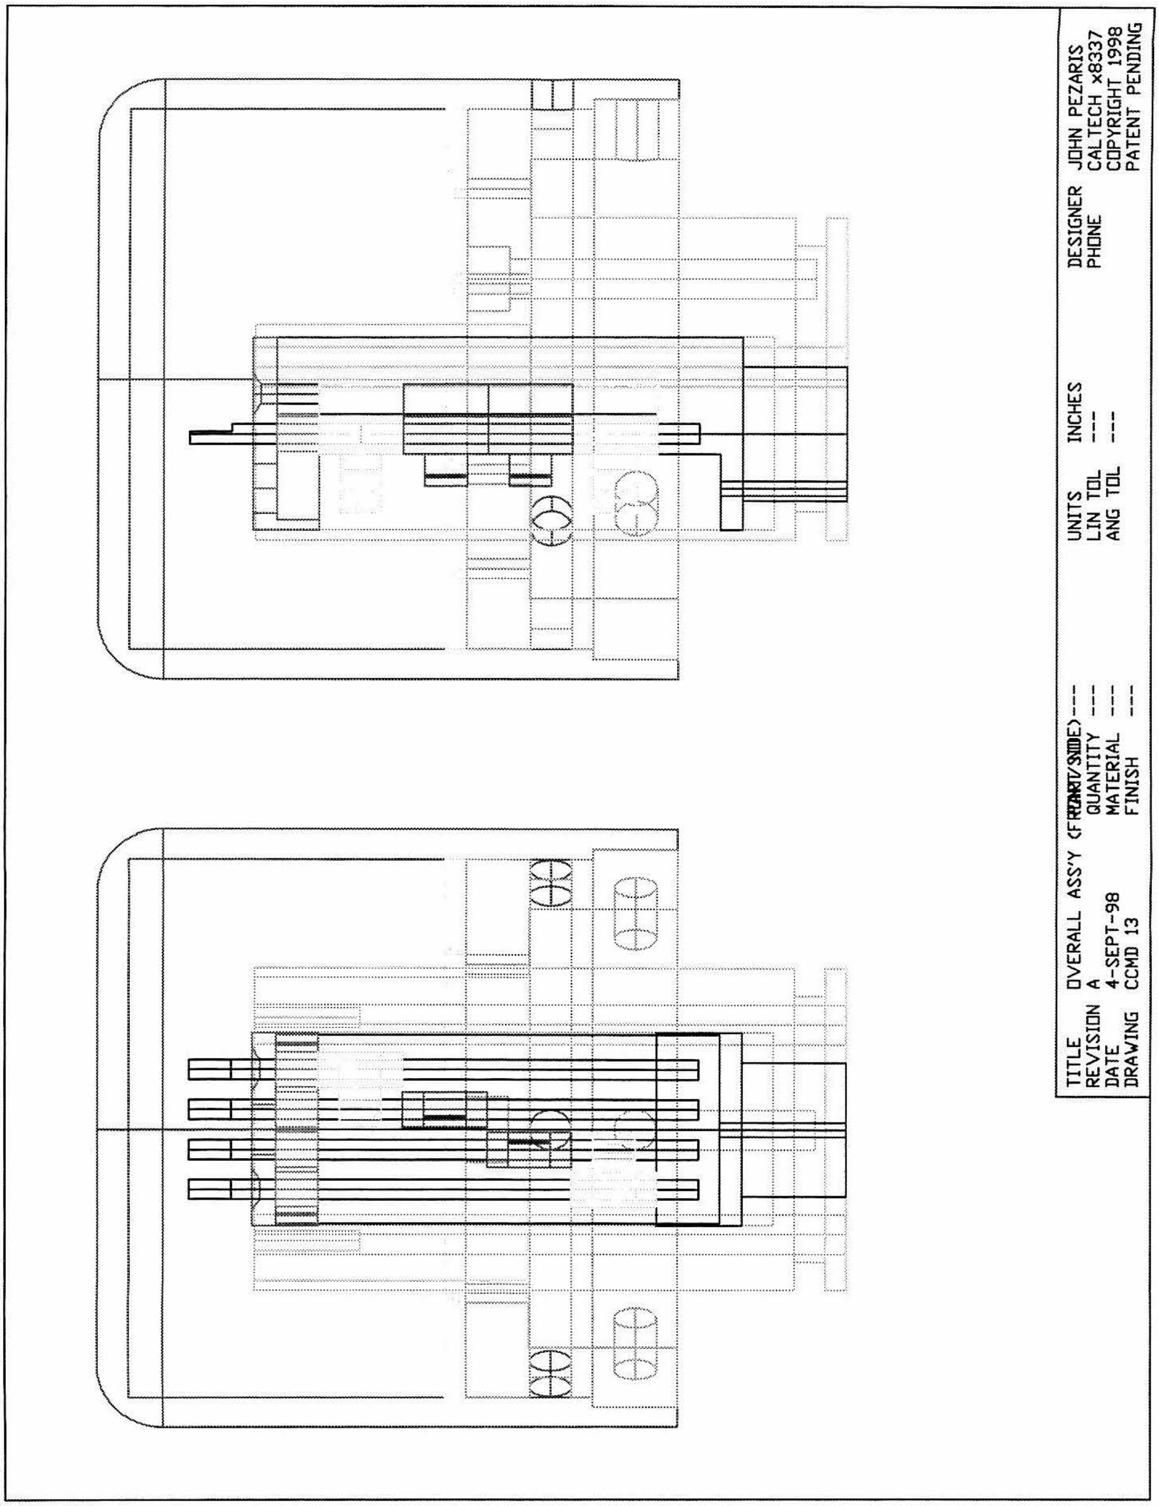

Figure 43: CCMD Overall Assembly

A wireframe diagram of the entire CCMD assembly, seen from two views.

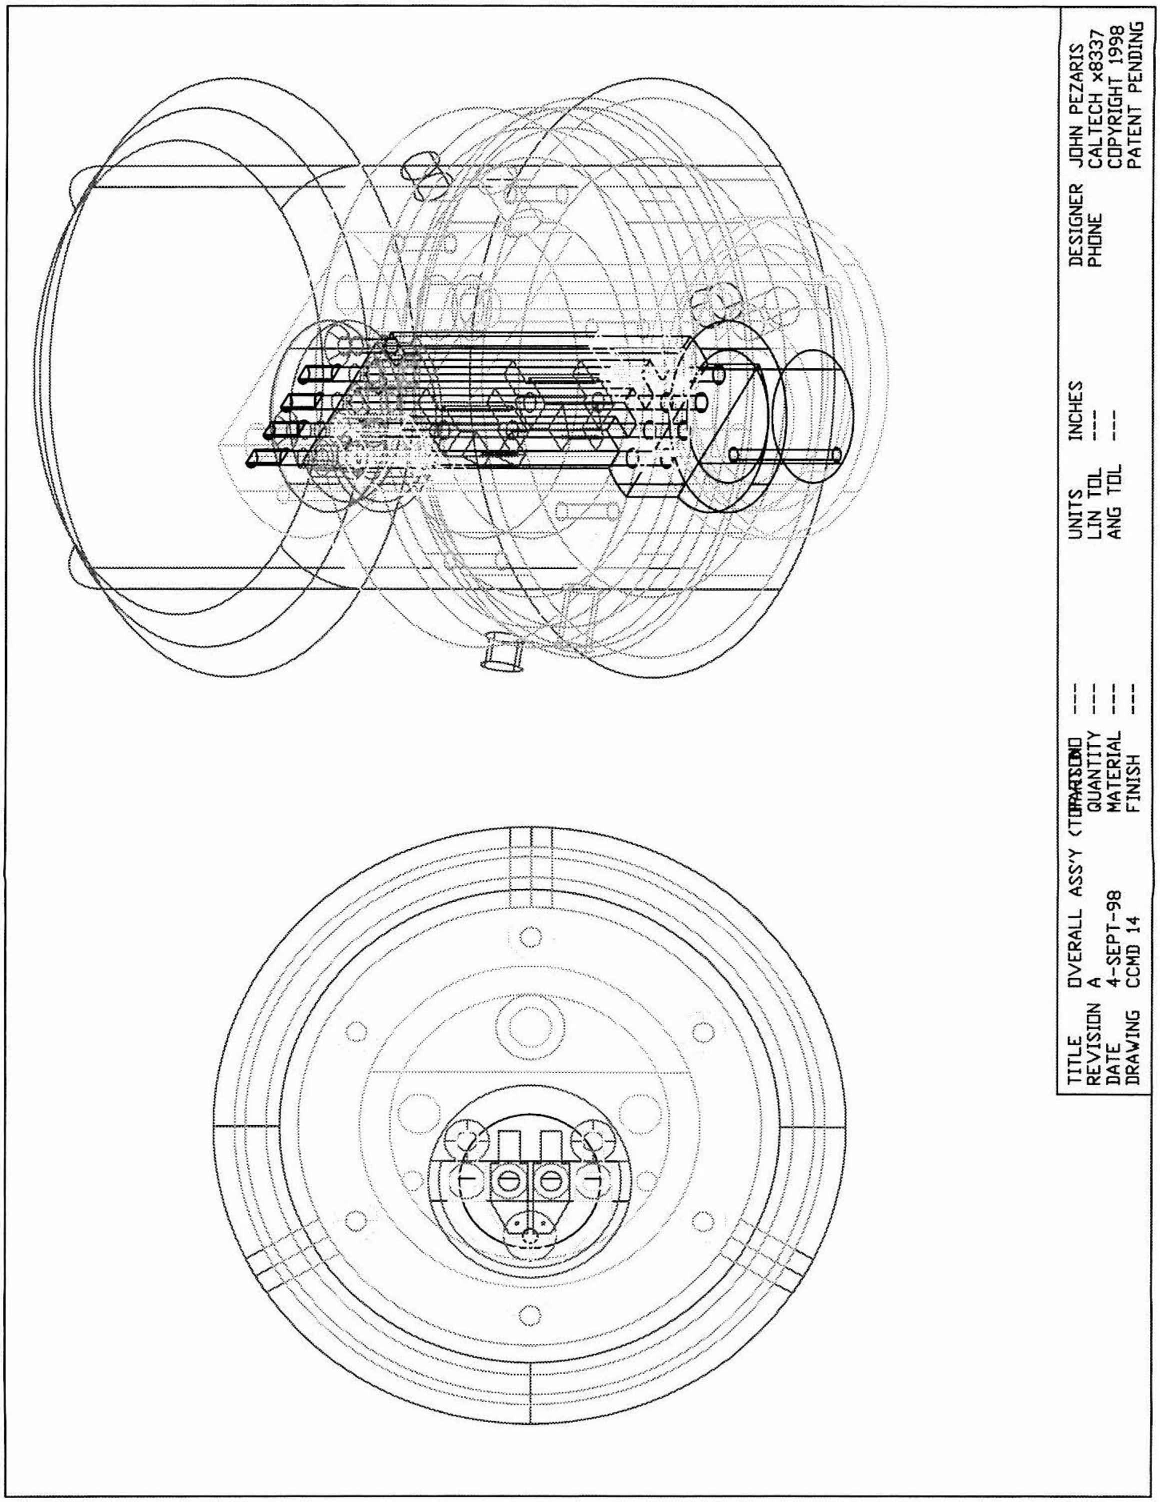

Figure 44: CCMD Overall Assembly
A wireframe diagram of the entire CCMD assembly, from another two views.

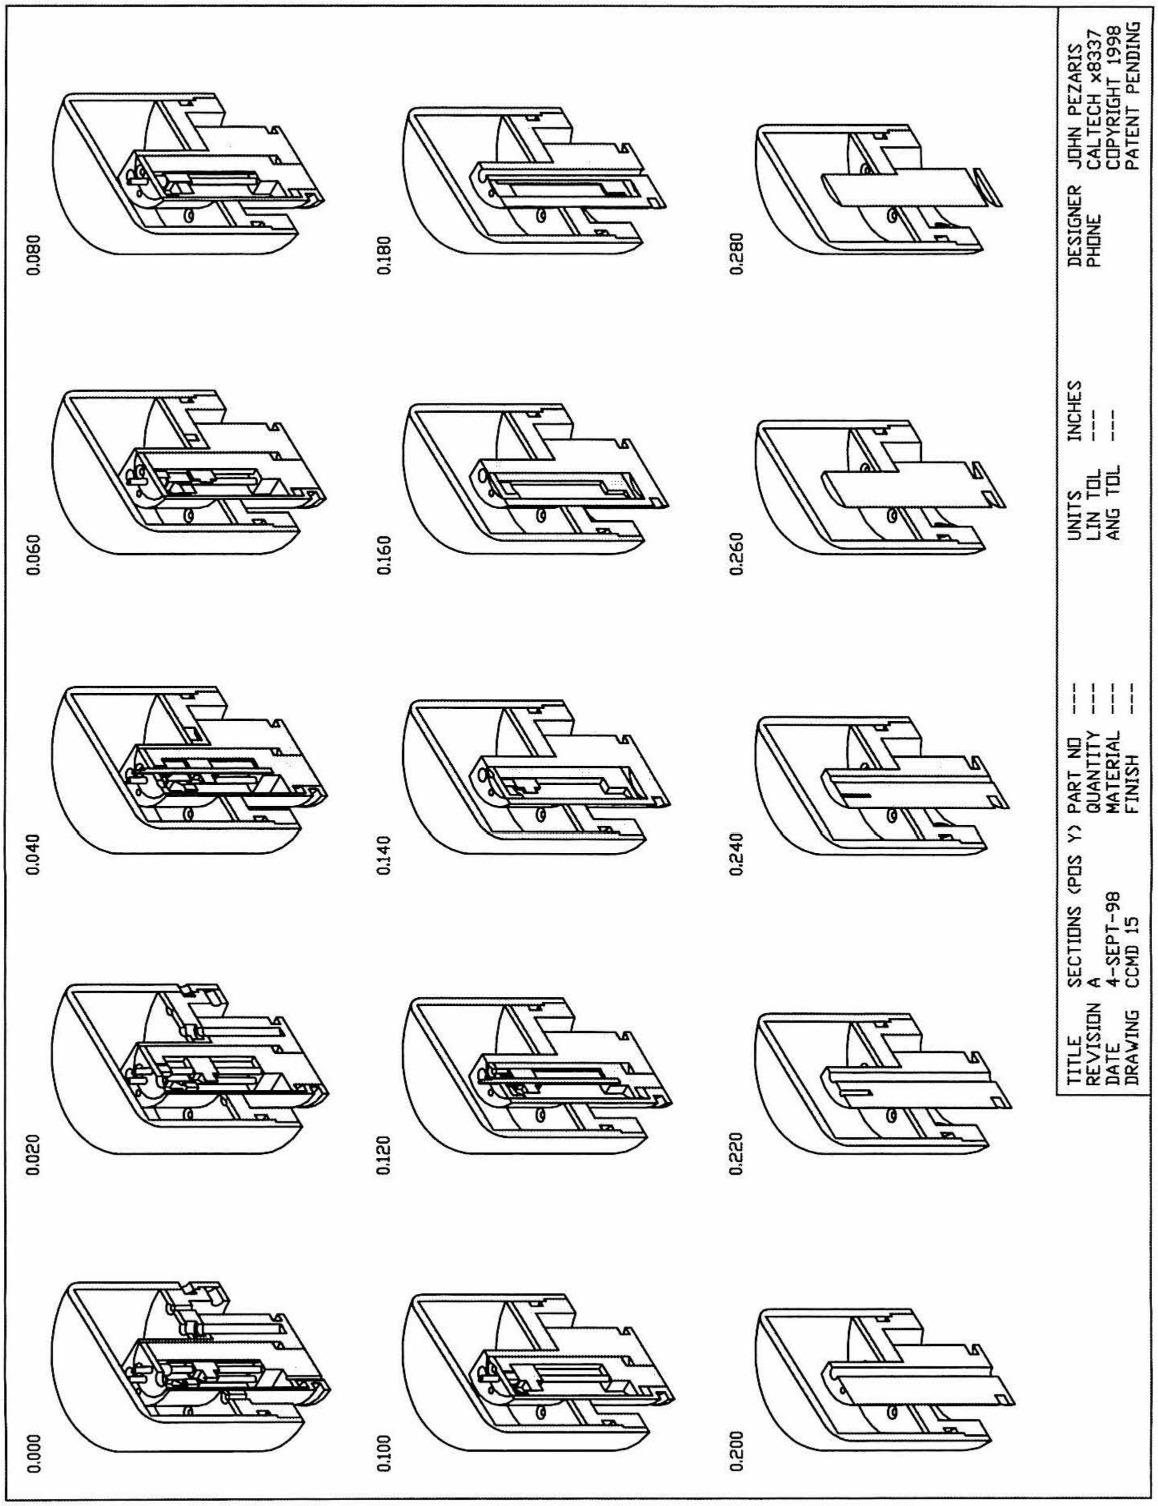

Figure 45: CCMD Sections (Positive Y)

Serial sections of the assembled CCMD in the positive y *direction. Sections are at 0.020 inch increments.*

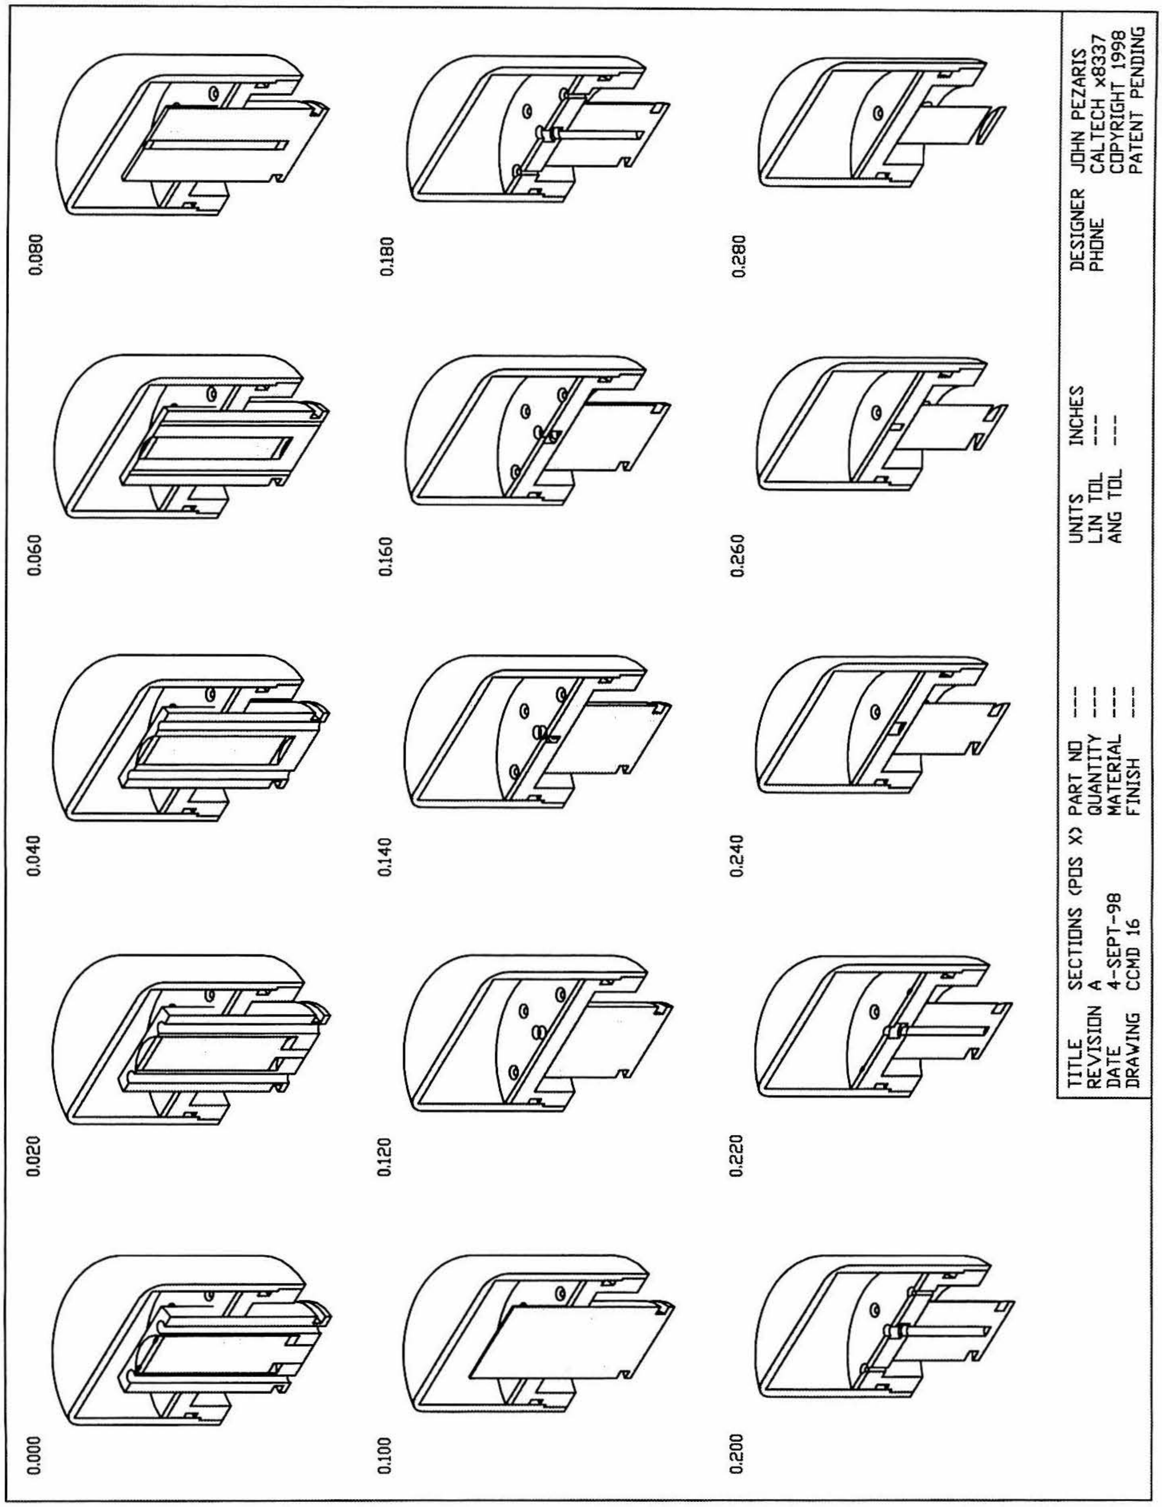

Figure 46: CCMD Sections (Positive X)
Serial sections of the assembled CCMD in the positive x direction. Sections are at 0.020 inch increments.

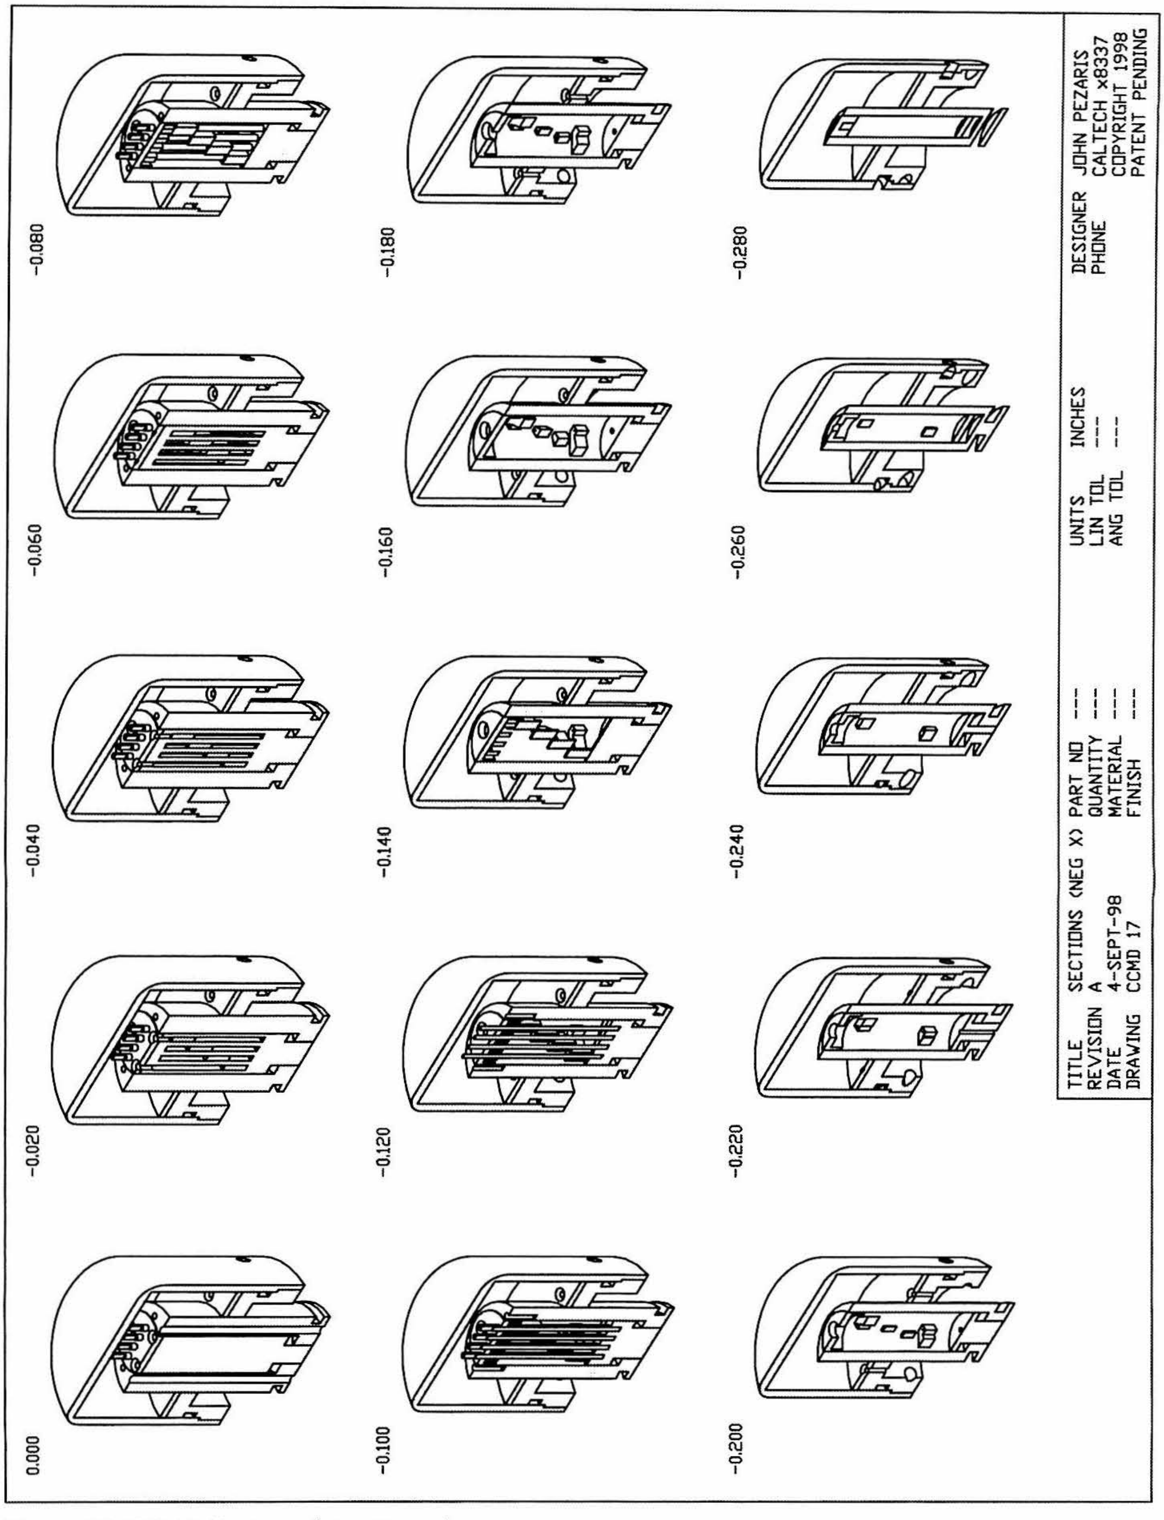

Figure 47: CCMD Sections (Negative X)

Serial sections of the assembled CCMD in the negative x direction. Sections are at 0.020 inch increments.

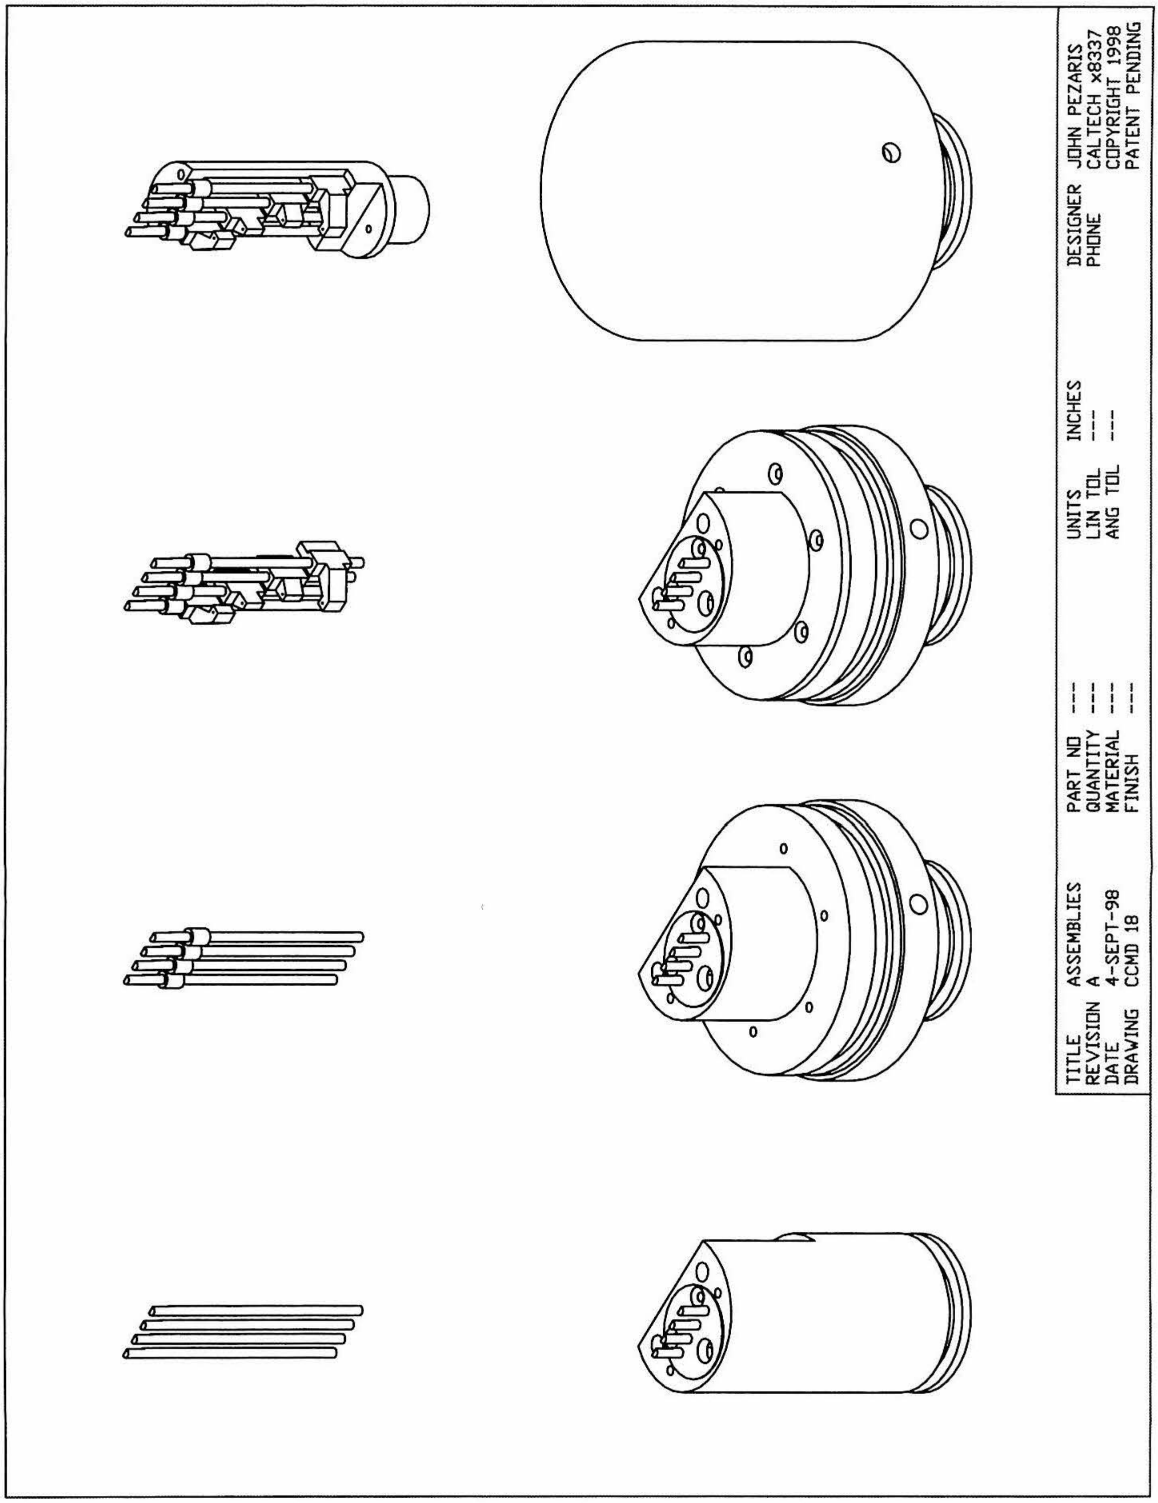

Figure 48: CCMD Assemblies/A
A series of figures showing the assembly of the CCMD in stages.

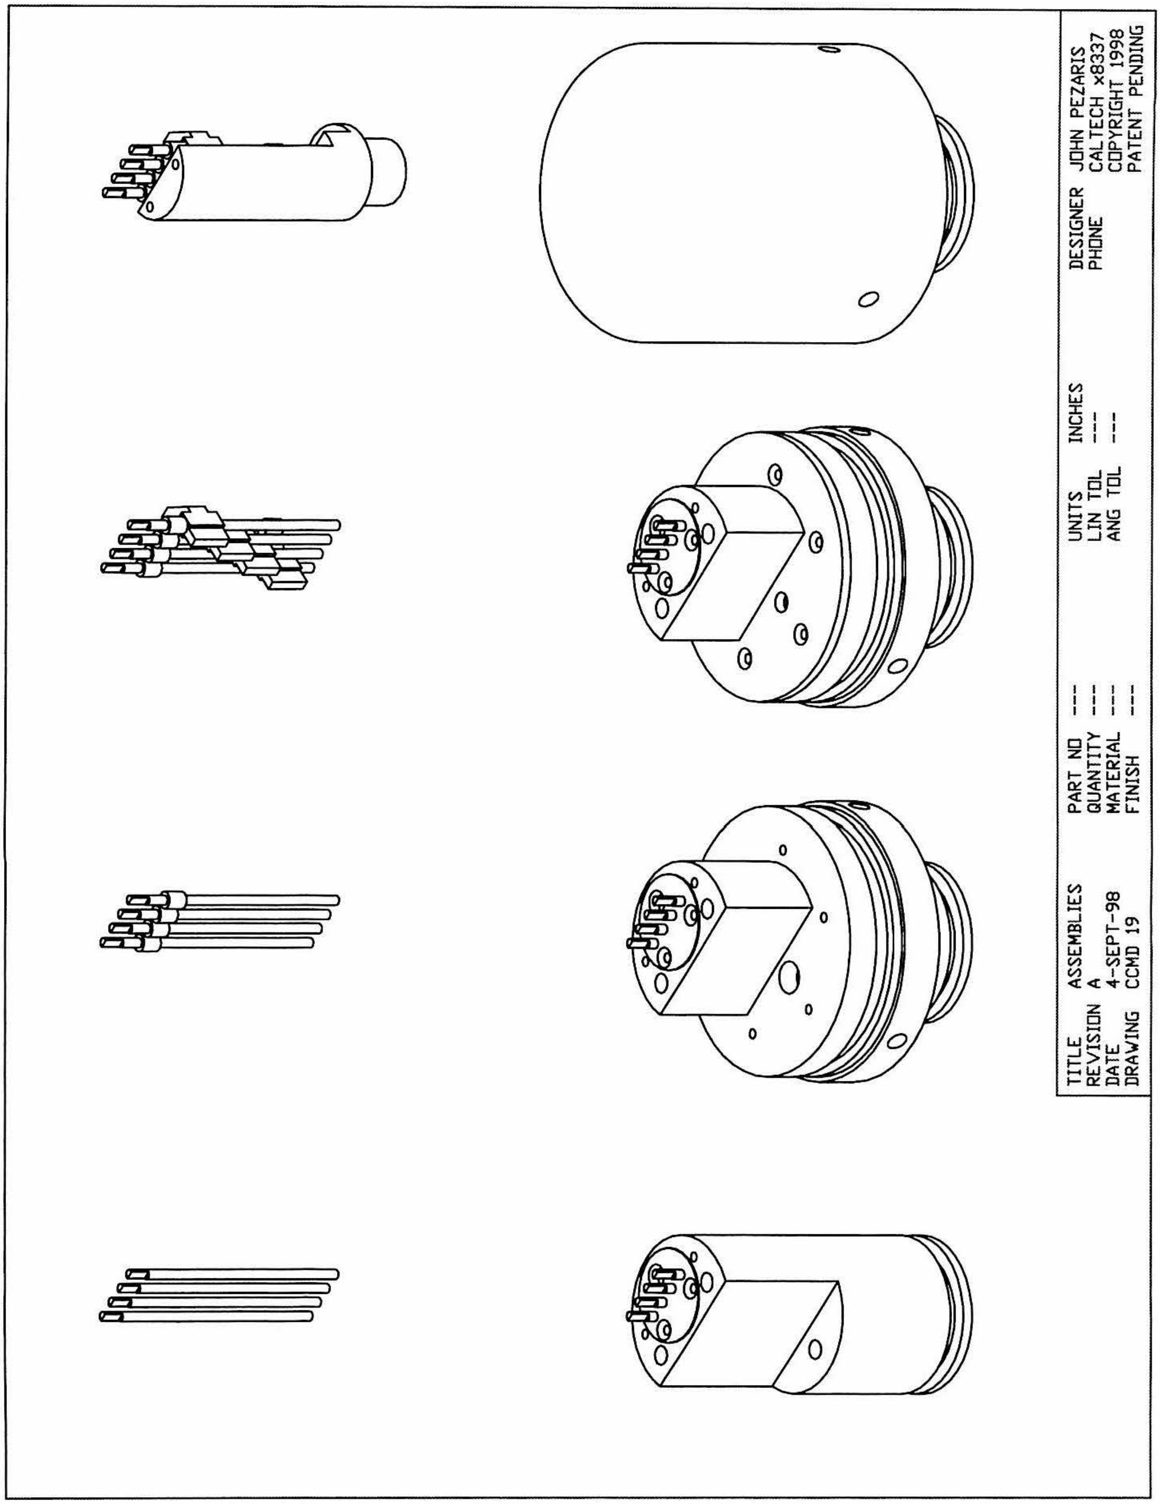

Figure 49: CCMD Assemblies/B
A series of figures showing the assembly of the CCMD in the same stages as the previous diagram, from the other side.

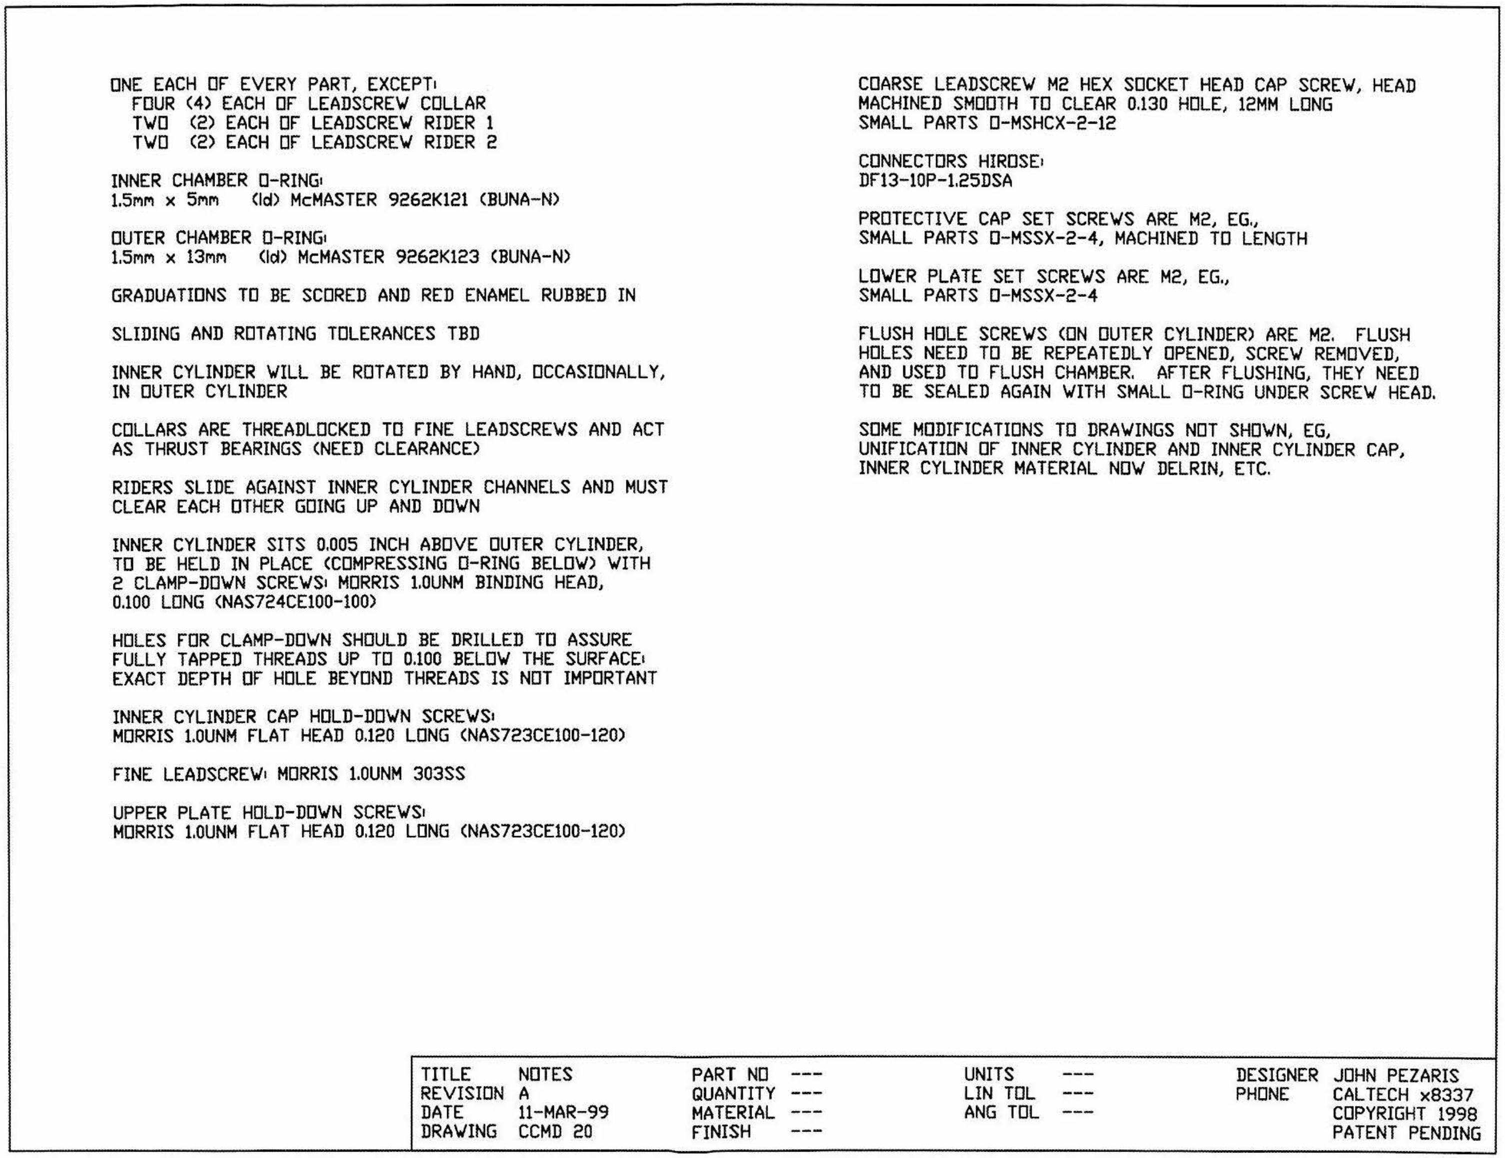

ONE EACH OF EVERY PART, EXCEPT:
  FOUR (4) EACH OF LEADSCREW COLLAR
  TWO (2) EACH OF LEADSCREW RIDER 1
  TWO (2) EACH OF LEADSCREW RIDER 2

INNER CHAMBER O-RING:
1.5mm x 5mm (Id) McMASTER 9262K121 (BUNA-N)

OUTER CHAMBER O-RING:
1.5mm x 13mm (Id) McMASTER 9262K123 (BUNA-N)

GRADUATIONS TO BE SCORED AND RED ENAMEL RUBBED IN

SLIDING AND ROTATING TOLERANCES TBD

INNER CYLINDER WILL BE ROTATED BY HAND, OCCASIONALLY, IN OUTER CYLINDER

COLLARS ARE THREADLOCKED TO FINE LEADSCREWS AND ACT AS THRUST BEARINGS (NEED CLEARANCE)

RIDERS SLIDE AGAINST INNER CYLINDER CHANNELS AND MUST CLEAR EACH OTHER GOING UP AND DOWN

INNER CYLINDER SITS 0.005 INCH ABOVE OUTER CYLINDER, TO BE HELD IN PLACE (COMPRESSING O-RING BELOW) WITH 2 CLAMP-DOWN SCREWS: MORRIS 1.0UNM BINDING HEAD, 0.100 LONG (NAS724CE100-100)

HOLES FOR CLAMP-DOWN SHOULD BE DRILLED TO ASSURE FULLY TAPPED THREADS UP TO 0.100 BELOW THE SURFACE: EXACT DEPTH OF HOLE BEYOND THREADS IS NOT IMPORTANT

INNER CYLINDER CAP HOLD-DOWN SCREWS:
MORRIS 1.0UNM FLAT HEAD 0.120 LONG (NAS723CE100-120)

FINE LEADSCREW: MORRIS 1.0UNM 303SS

UPPER PLATE HOLD-DOWN SCREWS:
MORRIS 1.0UNM FLAT HEAD 0.120 LONG (NAS723CE100-120)

COARSE LEADSCREW M2 HEX SOCKET HEAD CAP SCREW, HEAD MACHINED SMOOTH TO CLEAR 0.130 HOLE, 12MM LONG
SMALL PARTS O-MSHCX-2-12

CONNECTORS HIROSE:
DF13-10P-1.25DSA

PROTECTIVE CAP SET SCREWS ARE M2, EG.,
SMALL PARTS O-MSSX-2-4, MACHINED TO LENGTH

LOWER PLATE SET SCREWS ARE M2, EG.,
SMALL PARTS O-MSSX-2-4

FLUSH HOLE SCREWS (ON OUTER CYLINDER) ARE M2. FLUSH HOLES NEED TO BE REPEATEDLY OPENED, SCREW REMOVED, AND USED TO FLUSH CHAMBER. AFTER FLUSHING, THEY NEED TO BE SEALED AGAIN WITH SMALL O-RING UNDER SCREW HEAD.

SOME MODIFICATIONS TO DRAWINGS NOT SHOWN, EG, UNIFICATION OF INNER CYLINDER AND INNER CYLINDER CAP, INNER CYLINDER MATERIAL NOW DELRIN, ETC.

TITLE	NOTES	PART NO	---	UNITS	---	DESIGNER	JOHN PEZARIS
REVISION	A	QUANTITY	---	LIN TOL	---	PHONE	CALTECH x8337
DATE	11-MAR-99	MATERIAL	---	ANG TOL	---		COPYRIGHT 1998
DRAWING	CCMD 20	FINISH	---				PATENT PENDING

Figure 50: CCMD Notes
Drawing notes.

It looks like one of our thermopods.
But it's a very bad design.
— JOHN PARKER (Buckaroo Banzai, 1984)

Chapter Five: Instrumentation

1 Introduction

As this project was developing, it quickly became clear that quality of instrumentation was to be of paramount importance. We were striving to understand and classify signals which were traditionally considered noise. It brings us pleasure to see that the field is heading to the same position that we have been promulgating for these past five years or so: It is now not unusual to hear of experimental setups which capture multiple continuous voltage streams at 50 kHz sampling rates, a level of performance previously unheard of, and not necessary unless advanced classification algorithms are available. In fact, using 16-bit analog-to-digital (A/D) converters instead of more conventional 12-bit converters allowed us to carefully characterize the noise in our system. Doing this proved a significant turning point in the development of our spike sorter, as will be seen in a later chapter.

At the time we started the project, there were no readily available four-channel headstage amplifiers which met our various stringent criteria. We wanted to create a system where the unshielded high-impedance wire was kept to an absolute minimum, input bias current was similarly minimized, but most importantly, input current noise was as low as possible, thus devices such as sold by Fred Haer Corp, Alpha and Omega, RC Electronics, and the like, were inadequate. Accordingly, we set about to design our own.

At the suggestion of colleagues in the Konishi laboratory at Caltech, we decided to base our instrumentation around equipment produced by Tucker Davis Technologies (TDT). Their excellent line of instrumentation-grade computer-controlled componentry, including anti-alias filters, A/D converters, and the like, electrically decouples the critical analog stages from the computer through an optical fiber interface. Thus we bought the appropriate modules from TDT and built a few of our own to support tetrode recording.

We designed and constructed three items: (1) a four-channel low-noise headstage amplifier, (2) a four-channel switchable gain secondary amplifier, and (3) a quad double-throw single-pole electrically controllable switch. Various prototypes were built and tested, including the final equipment used in the project. Some early versions continue to see duty with other researchers in our lab. The following sections will discuss in detail the headstage and secondary amplifiers and switch, and are based in large part on previously written technical reports. But first, an overview of the experimental setup to place these designs in an operational context. Parts of this description will be revisited in a later chapter when the experimental methods are described.

2 Data Acquisition and Behavioral Control

In Figure 51, a block diagram of the data acquisition and behavioral control system is presented. The complexity of this diagram speaks to the level of difficulty when working with the awake behaving preparation in a high-fidelity recording experiment. A photograph of the actual setup can be seen in Figure 52, with various parts of the equipment labelled (although since not everything is visible, the list is necessarily incomplete). The system can be broken down into two broad sections, neuroelectrical and behavioral, which correspond to the upper and lower halves of the diagram, respectively. Additionally, there is a physical separation between the preparation which is placed in a shielded experimental chamber, and the control apparatus placed in an equipment rack just outside. This division corresponds, roughly, to the left and right parts of the diagram, split at the first third. However, we will take the neuroelectrical-behavioral split as primary and describe the two halves in turn.

The neuroelectrical path consists, as depicted, as a pair of paths, one for each of two tetrodes. Each of the two paths is identical up until they join in the data capture computer. The path starts with a tetrode mounted in a microdrive; the microdrive position is set via hydraulic line from a controller in the experimental rack, and once adjusted correctly, the electrode tip will be in an experimentally relevant part of the subject's brain. Signals transduced by the tetrode are amplified by a low-noise headstage, THA2, which includes a modicum of DC decoupling and low-pass filtering. The output of the headstage travels along a long cable to outside the experimental chamber to an equipment rack. Signals are then additionally amplified at AMP4 to an adjusted level (nominally $A = 200$), anti-alias filtered at 10 kHz, and digitized at 24–50 kHz. The A/D converters are synchronized through an external clock. One tetrode channel is split out of the four and routed through a single pole double throw switch and back; this allows the impression of an out-of-band signal, such as behavioral timing marks, onto the neural stream. The A/D output is sent over an optical fiber link to a DSP card in the data capture computer which feeds archiving and analysis software. The four amplified signals are also available for viewing on an oscilloscope, monitoring over headphones, or traditional window discrimination. In typical use, two channels of the four are selected for audio monitor, and one for window discrimination. Although not shown, there is an additional stage of filtering contained within the window discriminator, nominally set as a 30–6000 Hz bandpass. The window discriminator allows on-line monitoring, and the output is sent for analysis and display to the behavioral control computer.

Despite the two-tetrode presentation in Figure 51, the actual system has the capacity for four. To generalize the diagram, the pair of tetrode paths as shown would be replicated so that two tetrodes are fed into a single computer. Additionally, the actual system has a single shared oscilloscope for tetrode monitoring, only a single window discriminator, and an oscilloscope for general purpose use including monitoring eye position or the output of the window discriminator.

The behavioral path consists of means to control the experimental environment and to monitor the animal's behavior. The experimental environment includes room lights, juice reward, as well as the optical stimuli, and microstimulation. The behavioral monitoring includes an infrared video image of the animal over closed-circuit TV and measurement of the animal's eye position through an implanted scleral search coil.

The various system outputs, room lights, reward, etc., are generated through TTL and analog outputs which can be seen in the center of the lower half of Figure 51. Following this column from top to bottom, we see one TTL signal being sent to the JUICER which, when activated, provides the animal with liquid reward. The computer-generated signal can also be augmented with a manual input, seen to the left of the TTL I/O box. Continuing downwards, we have another TTL line controlling the room lights through a high-power 120V DC supply, followed by three TTL lines controlling the optical bench shutters. These shutters turn on and off the beams of light which generate the stimuli; the optical bench will be described in detail in a subsequent chapter. Below this is an eight-channel D/A converted which is used to generate three pairs of x-y signals to the three galvo-driven mirrors which position the three independent optical stimuli. Finally, one channel of the D/A output is used to drive a microstimulator for exploratory work, although this device was not used in the current research.

The behavioral monitoring consists of an eye position measuring system for fine control of the experiment, augmented with an infrared video image for broader-scale feedback to the experimenter of the animal's state. The eye position system uses a variant of the standard search coil technique [30, 63] which consists of measuring the mutual inductance between large horizontal and vertical primary coils and a small, shared secondary coil, in an air-core transformer where the secondary coil is free to rotate about the vertical and horizontal axes. For small deviations from a central position, the mutual inductance, as measured by the voltage coupled to the eye coil through high-frequency magnetic fields, varies approximately linearly with angular deviation. Thus the EYE POS box generates two synchronized fields at rationally related frequencies, one which is driven to a pair of coils above and below the subject to measure vertical deflection, and the other driven to a pair of coils to the left and right of the subject for horizontal deflection. Decoding of the mixed HF signal from the eye coil is done by synchronous switching, producing independent analog values for x and y deviation. The coils can be seen to the left in the diagram of Figure 51, and on the right in the photograph of Figure 52.

Now that the experimental setup has been reviewed, providing context for additional detail, the remainder of this chapter will cover three of the subsystems designed and constructed for this project: the headstage amplifier THA2, the secondary amplifier AMP4, and a switching module SW1.

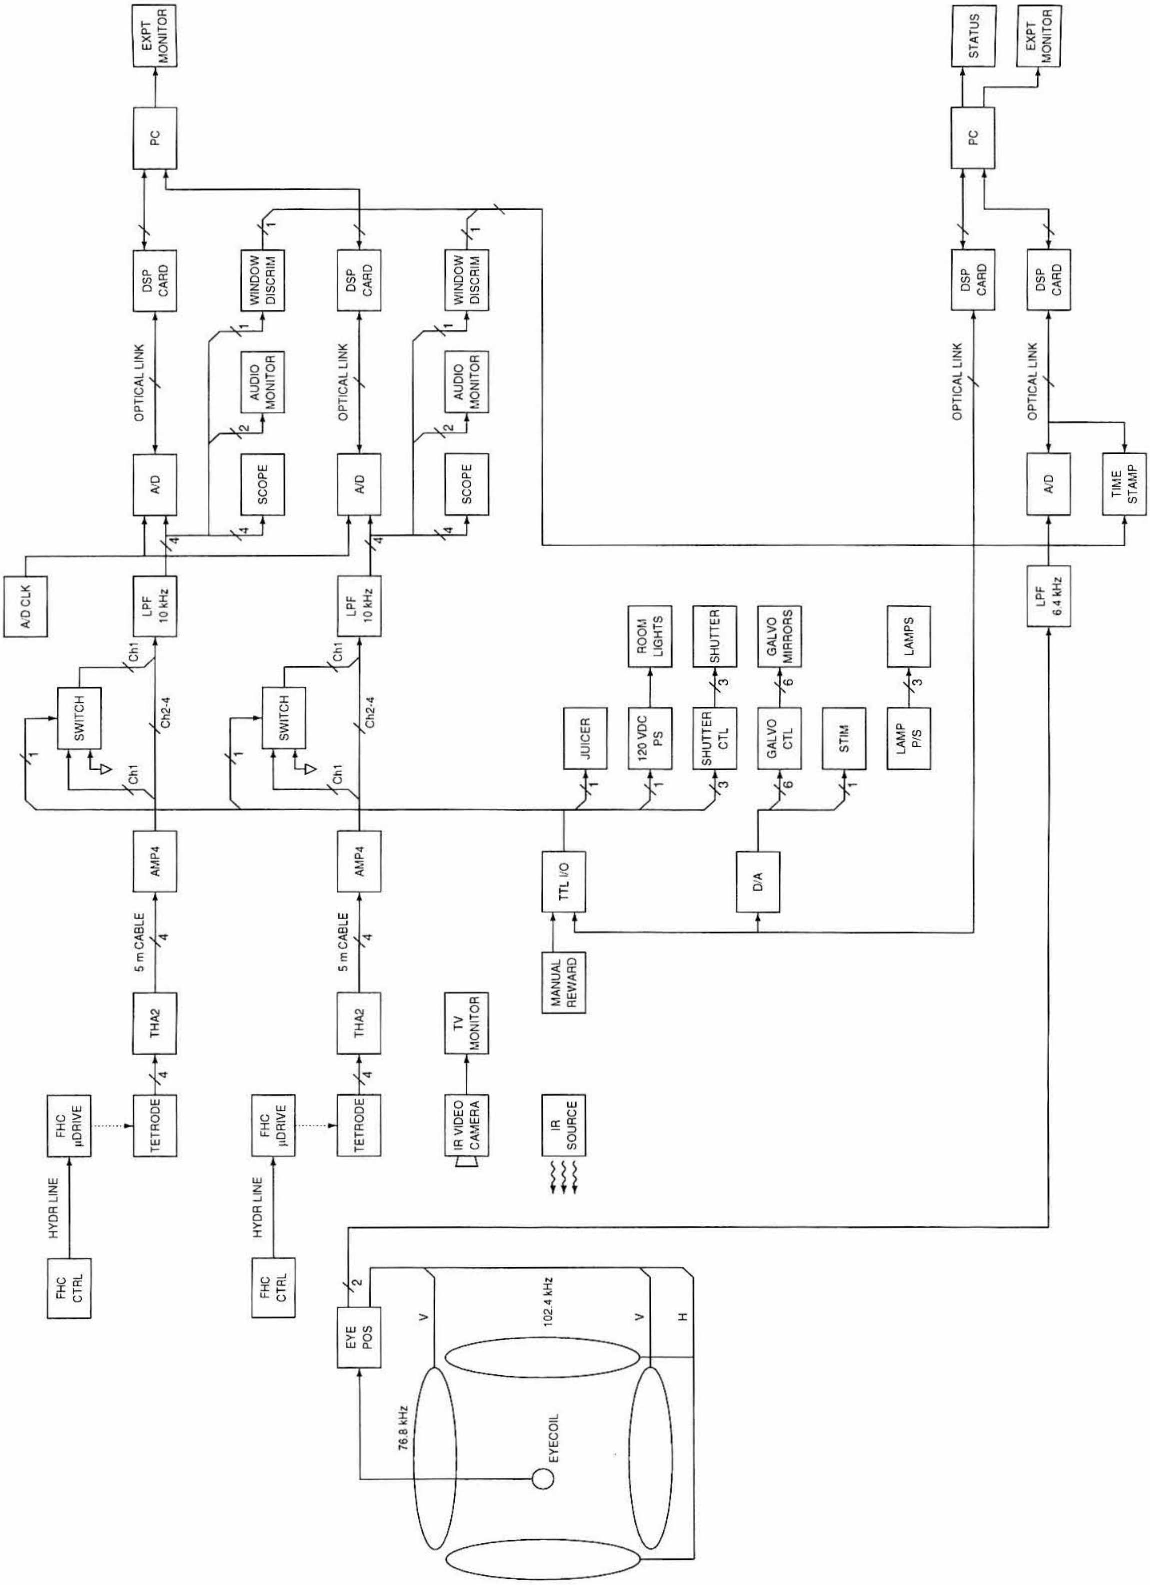

Figure 51: System Diagram

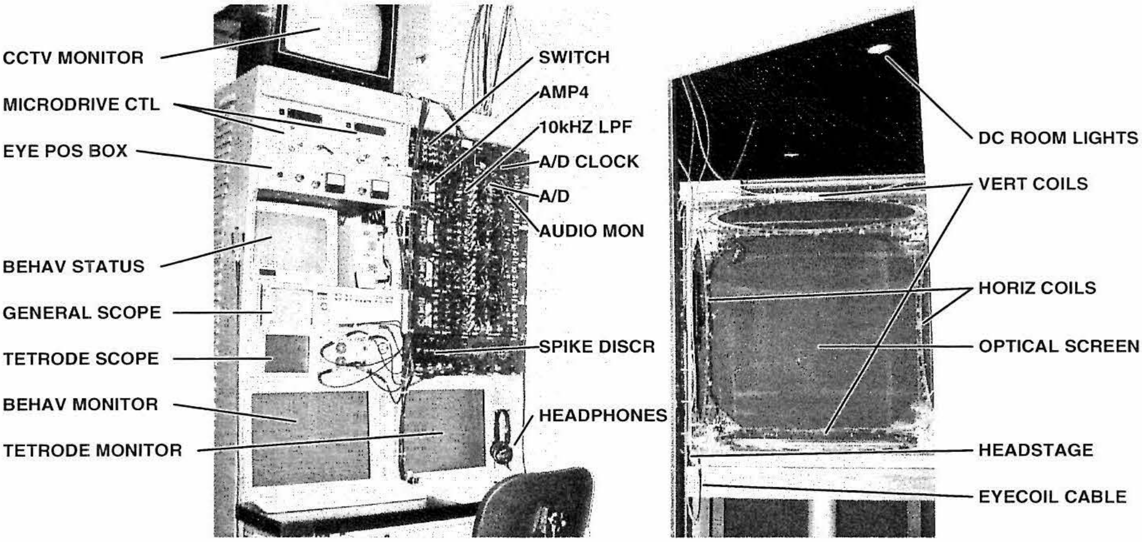

Figure 52: System Photograph

3 A Miniature Four Channel Headstage Amplifier

The first stage of amplification in an electrophysiology system is the most important. This section will present the design and construction of a miniature four-channel headstage amplifier for use with tetrodes, including specification, response characteristics, schematics, printed circuit board layout, parts lists, and suppliers. Our results indicate that for frequencies above 10 Hz, and electrode impedances above 100 kΩ, the electrical noise floor presented by this design is below the anticipated thermal noise of the electrode.

3.1 Introduction

Proper signal conditioning when making neural recordings requires a first stage of amplification and impedance transformation that is located physically close to the recording electrode. This requirement of proximity is primarily due to the relatively high source impedance of electrodes typically used for neural electrophysiology, as even short signal runs become subject to pickup from external sources. In turn, proximity requires miniaturization of the electronics, and when multi-channel electrodes are used, such as tetrodes, this requirement is compounded.

One standard solution to this problem is to use a single stage of impedance matching with a lone JFET follower. The advantages to this approach are low-power and extremely small size, while the disadvantages are higher noise levels, lack of gain and filtering, and lack of sufficient drive to power substantial cable runs. Trying to solve all of the design constraints with one transistor means compromising on nearly every aspect of circuit performance.

Therefore, the approach taken in this project was to use commercially available operational amplifiers in surface mount packages. While this has the disadvantage of being substantially larger than some approaches (especially since, as will be discussed, components selected were far from the most compact available), it is still quite small, and has the advantage of being able to optimize input, transfer, and output characteristics. The result, as presented here, is a very low-noise, low-power, miniature four channel AC amplifier with second-order high-pass rolloff, that is robust and can drive a long cable with ease.

The design uses a small printed circuit board which holds all of the circuitry, directly connects on one end to a recording electrode, and at the other through a 5 meter cable to a secondary amplifier and power source. Without the cable and associated connector, the headstage measures approximately 15 mm by 60 mm and weighs some 10 grams. It accepts signals from high impedance electrodes, amplifies with a gain of 100, and filters with a highpass rolloff of 0.1 Hz and a low-pass rolloff of 15 kHz.

3.2 Schematic and Design Description

The schematic for the tetrode headstage amplifier (THA2) can be seen in Figure 53. Signals arriving at the input connector are filtered with a DC blocking capacitor and a bias-current bleed-off resistor forming a first-order high-pass filter, amplified by the first stage which has a first-order low-pass rolloff, and further amplified by the second stage, also with a first-order low-pass rolloff, before being driven into a long cable.

Because the source impedance from neural electrodes is typically 0.1–3 MΩ, and recorded voltages are on the order of tens to hundreds of microvolts, the current noise of the first stage of amplification is more important than its voltage noise. Further, to insure that bias currents from the amplifier do not affect the experimental subject, and to eliminate the wet cell voltage created by the electrode in tissue, a DC blocking capacitor is necessary, along with a bleed resistor to prevent bias currents from floating the input node into saturation. These two components are selected so that they form a high-pass filter with $f_c = 0.1$ Hz, a sufficiently low cutoff to retain signals of interest while providing the required isolation. The operational amplifier selected for this stage is National Semiconductor's LMC6082AIM; it has outstanding noise characteristics, low input bias current, good input offset voltage, reasonable power dissipation, and a reasonable gain-bandwidth product. It is used in a non-inverting configuration with a gain of 10, with a single low-pass zero inserted in the feedback path to reduce the gain to 1 for signals above 15 kHz. This feedback network also insures stability which is often a concern with these and similar low-noise, low-input current, fast amplifiers.

The second stage is optimized for driving cables, and thus, a higher output-current amplifier was chosen, and designed again in a non-inverting configuration with a single low-pass zero in the feedback path, reducing the nominal gain of 10 to a gain of 1 above 15 kHz. The operational amplifier selected for this stage is Linear Technology's LT1355, a high output-current, wide-bandwidth, low-noise amplifier that consumes a reasonable amount of power.

To insure high-quality isolation, power is delivered to the headstage at ±8V and locally regulated to ±5V. The incoming power lines are decoupled via 4.7μF tantalum capacitors before regulation through an LM78L05 for the positive supply and an LM79L05 for the negative supply. Each integrated circuit has local decoupling capacitors for both supplies, and the layout insures a low-impedance ground path. Local power regulation is important to prevent signals picked up in the power-supply leads from appearing in the opamp outputs; although the low-frequency power supply rejection of modern opamps is excellent, above 10 kHz or so, it can be as low as 20–40 dB.

The cable is driven from the second stage in a single-ended configuration, however, the cable has individual coax cables for each channel, and an overall shield. The signal ground is taken from the power supply ground and is kept separate from the overall shield. At the receiving end (not described in this document), the signal is differentially amplified for a pseudo-differential configuration, combining the better noise rejection of differential signaling, with the lower component count of single-ended signaling.

The overall shield is not made available at the electrode end of the headstage, but is brought out to the connector housing at the secondary amplifier end of the cable. At the secondary amplifier, it is connected to chassis ground. The combination of local power regulation, pseudo-differential signaling, and outer shield knocks external interference down to manageable levels, as presented below.

The primary use for these amplifiers is with four-wire bundle electrodes called tetrodes [60, 56]. The input connector was selected to provide a direct connection to the previously developed carrier tubes, and is a four-position header with gold-plated contacts. The output connector is a 12-pin locking circular connector which, although somewhat expensive, is small, rugged, reliable, and dense. The cable was selected to provide the desired layers shielding while also being flexible, so as to limit strain applied to the input connector and therefore the attached electrode, and fit the output connector.

3.3 PCB Layout

A two-sided printed circuit board layout, 0.450 by 1.975 inches (11.4 by 50.2 mm) when cut, is show in Figure 54. Each side of the board holds two channels, and signal flow is routed cleanly from the input connector (on the left edge in Figure 54) through to the output connector (on the right edge). Power is distributed through traces around the periphery, and ground along a central strip. Each longitudinal half of both sides is nearly symmetric, and the two sides are nearly identical, with the majority of the differences occurring at the local regulation near the cable soldering points. To minimize width and length and not require additional layers, a design was chosen which uses two jumpers, both at the local regulators, and both to transfer regulated power to the opposite side of the board.

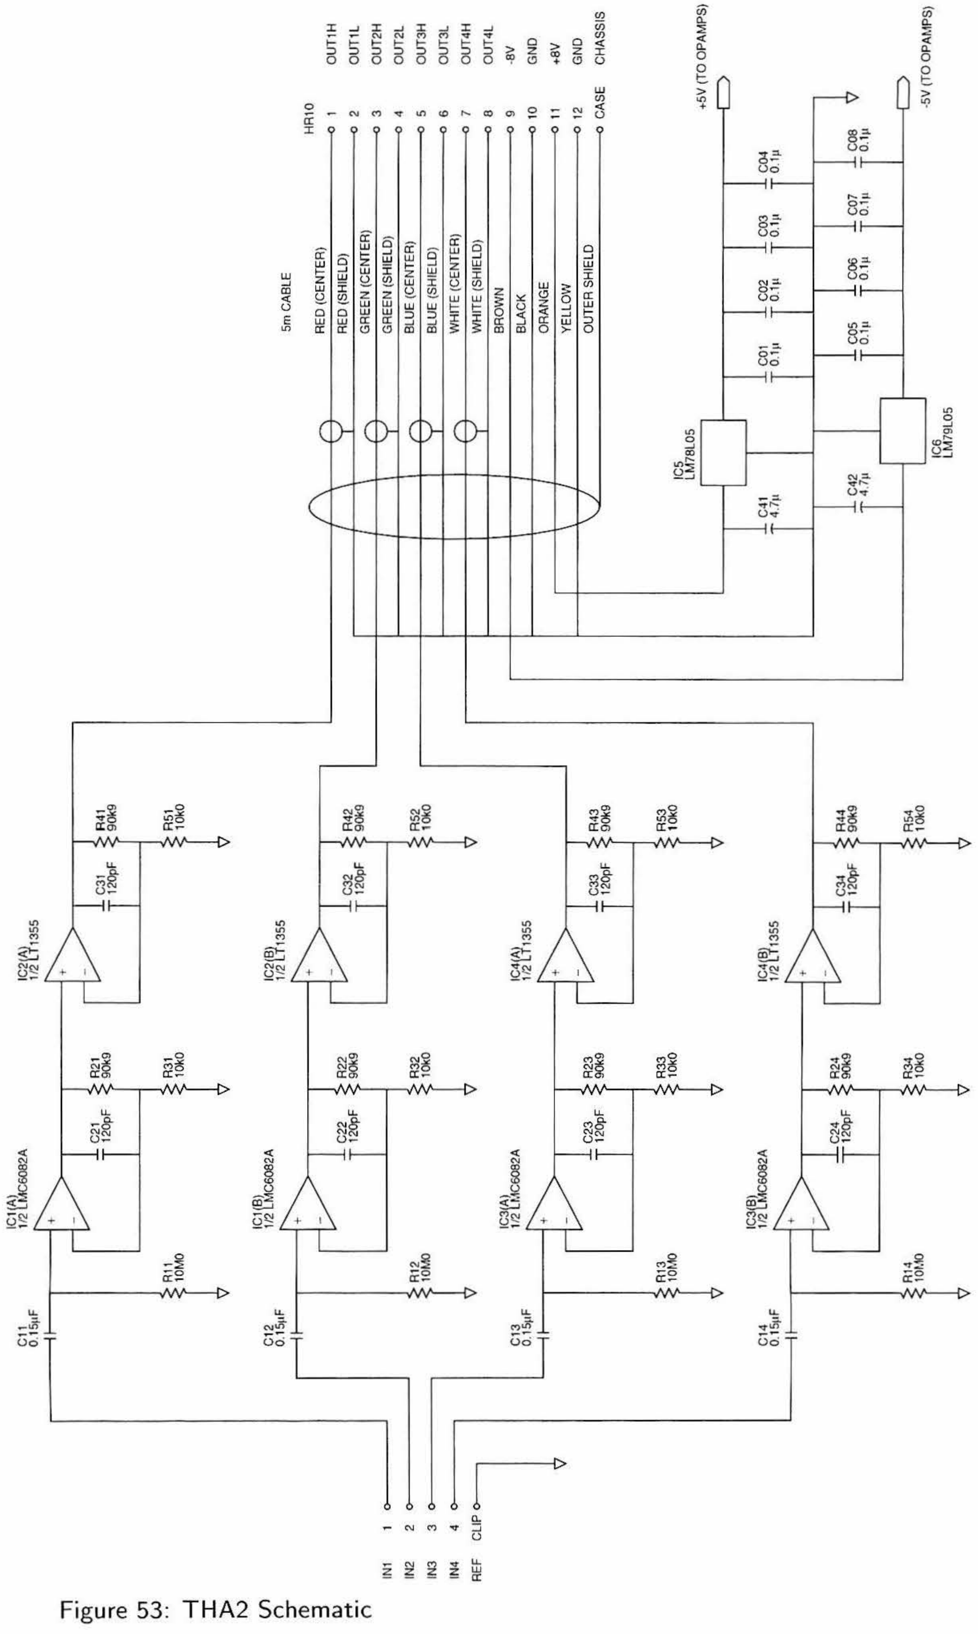

Figure 53: THA2 Schematic

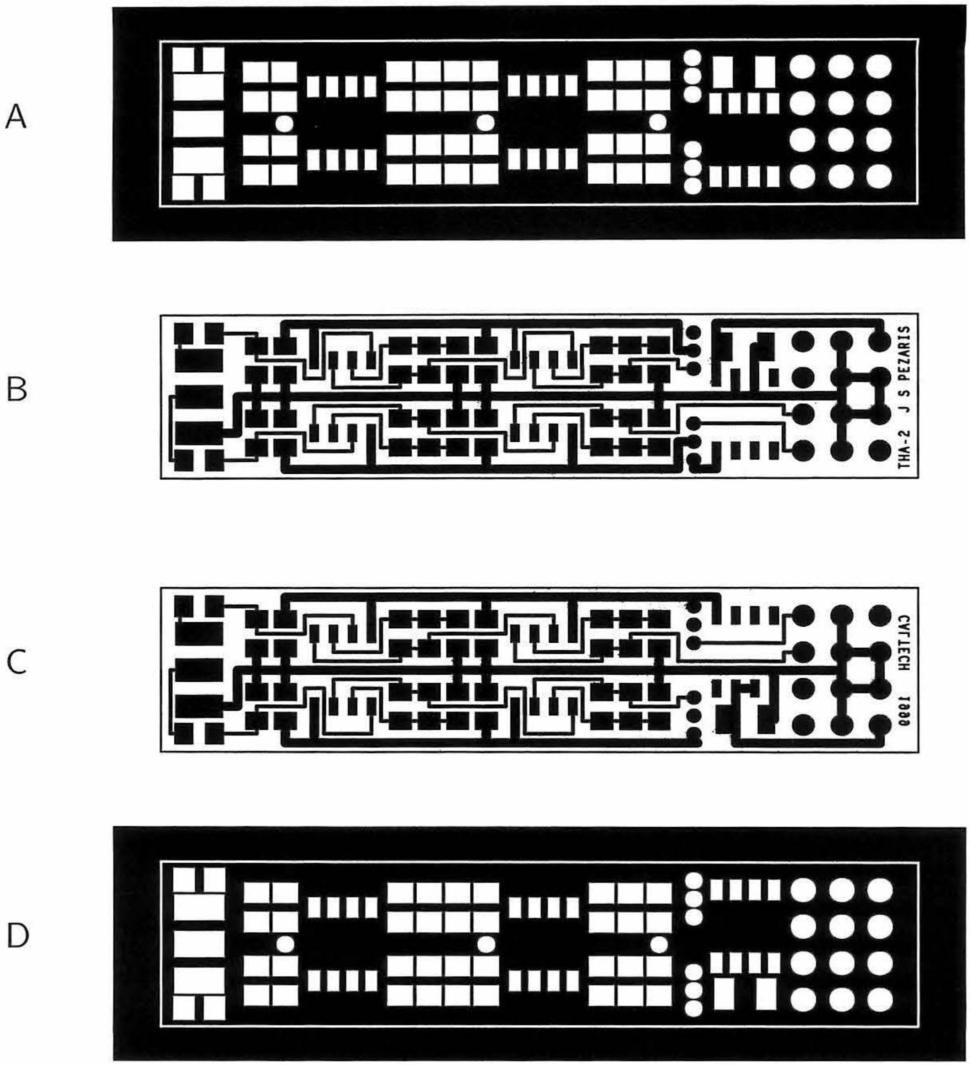

Figure 54: THA2 Circuit Board Layout
Layers used to construct printed circuit board, plotted at approximately twice life-size. Mask layers are presented in negative, and bottom layers in mirror image, as is standard for the industry. The rectangular outline represents the edge of the board, and actual size is 0.450 by 1.975 inches (approximately 11.4 by 50.2 mm). **A**, *top mask.* **B**, *top etch.* **C**, *bottom etch.* **D**, *bottom mask.*

3.4 Parts List

A list of required parts is given in Table 4. At the time of this writing, these parts are all readily available, and reasonably inexpensive. The author can be contacted for limited copies of the printed circuit board, or an excellent source for low-volume PCBs can be found in the sources section. Net lists and Gerber plots are available upon request for non-commercial use.

item	*part number*	*description*	*source*
PC1	THA1 (Custom)	printed circuit board	Advanced Circuits
IC1, IC3	LMC6082AIM†1	input stage opamp	Digi-Key
IC2, IC4	LT1355CS8†2	output stage opamp	Digi-Key
IC5	LM78L05ACM†2	5V positive regulator	Digi-Key
IC6	LM79L05ACM†2	5V negative regulator	Digi-Key
R11–R14	RK73H2AT1005F†3	0805 10M 1%	Surface Mount Dist.
R21–R24, R41–R44	P1005CCT†4	0805 10k0 1%	Digi-Key
R31–R34, R51–R54	P9095CCT†4	0805 90k9 1%	Digi-Key
C01–C08	ECJ-2YB1H104K†4	0805 0.10μF 50V X7R	Digi-Key
C11–C14	ECJ-2YB1E224K†4	0805 0.15μF 25V X7R	Digi-Key
C21–C24, C31–C34	ECU-V1H121JCG†4	0805 120pF 50V NPO	Digi-Key
C41–C42	ECS-T1CY457R†4	EIA(A) 4.7μF 16V tantalum	Digi-Key
W1	3M1181B-ND†5	shielding tape	Digi-Key
W2	EPS2034-1†5	shrink tube	Digi-Key
T1	929852-01-36†5	input connector (modified)	Digi-Key
T2	HR10-10P-12P†6	output connector	Digi-Key
X1	GC397-ND	reference input clip lead	Digi-Key
X2	RP101C-ND	cable tie	Digi-Key
X3	VW-4754-500BK†7	cable	Bi-Tronics

Table 4: THA2 Parts List
A complete list of parts for constructing one headstage amplifier. †1 National Semiconductor. †2 Linear Technology. †3 KOA. †4 Panasonic. †5 3M (Digi-Key part number). †6 Hirose. †7 Mogami (Bi-Tronics part number).

3.5 Construction

When manufactured, the printed circuit boards are paneled 5-up and routed with breakout points. This not only eases the job of the PCB manufacturer, it also eases handling during soldering.

Construction by hand, rather than by automated mechanism, requires a steady hand, good soldering skill, a fine-tipped soldering iron, and bountiful patience. Mount the lower-profile, passive components first, then follow with the active components, then the taller passive components, the input connector, and lastly the output cable and reference lead. When soldering each component, it helps to tin a minute amount of solder to one pad before placing the component and applying heat at that pad to fix the component's position. After that, the remaining leads can be soldered with relative ease. Various stages of assembly are depicted photographically below.

After all wiring is completed, the board is protected with a series of layers. The first is a single layer of black electrical tape for insulation, followed by a layer of carefully overlapped foil tape which is brought in contact with the cable's overall shield. Finally, a layer of heat shrink tube is applied to protect the foil and provide strain relief for the cable-to-board solder joints.

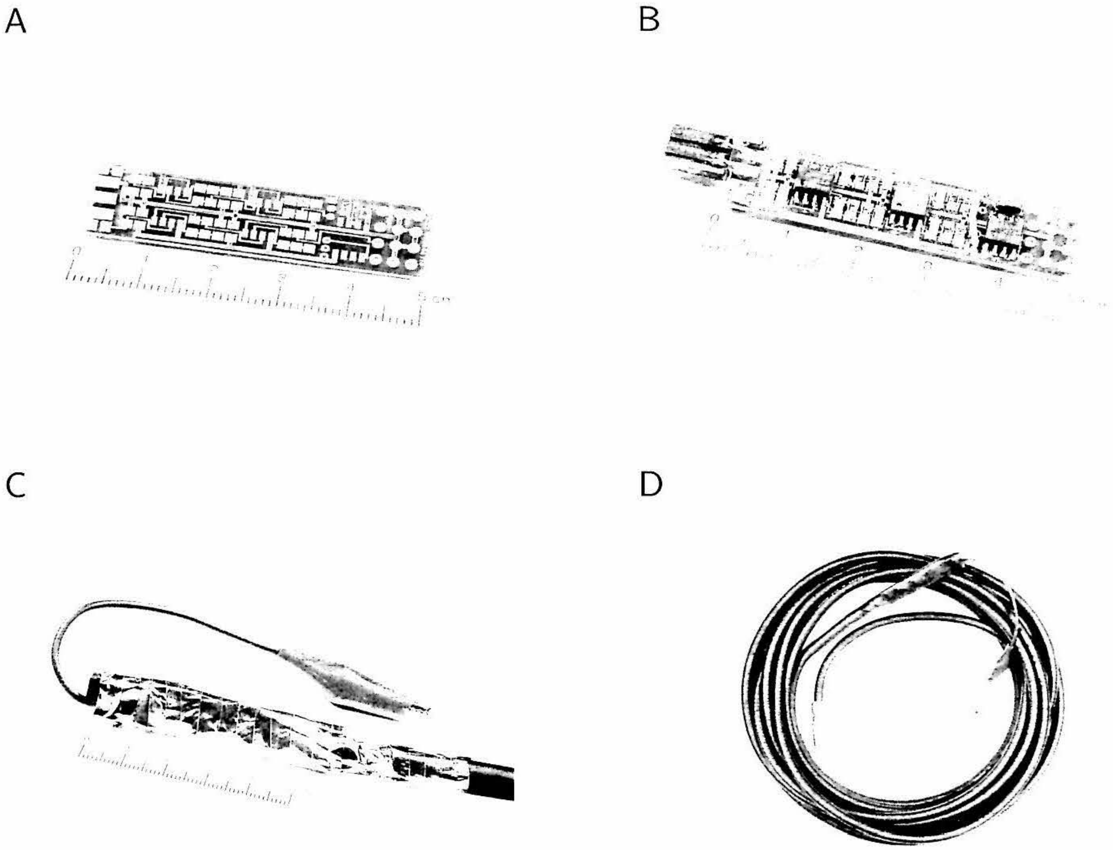

Figure 55: THA2 Assembly
Four views during assembly of the headstage. **A**, *a blank board as delivered from the PCB manufacturer.* **B**, *fully populated board.* **C**, *board with cable and ground lead attached, showing layer of copper shielding tape.* **D**, *completed headstage with shrink tube, cable, and connector.*

3.6 Operation

Operation is straightforward. The headstage is first connected to the secondary amplifier to provide a source of power, and then attached to the electrode and a ground reference. No adjustments are necessary or even possible. During normal operation, the headstage gets slightly warm to the touch.

3.7 Results

These headstages have been in use in our laboratory at Caltech for more than a year and have supported hundreds of recordings. They work well, are reliable, and produce excellent neural data. Data were collected for analysis below using the headstage in conjunction with a mating secondary amplifier which provides additional gain, and additional filtering.

A sample neural recording is presented in Figure 56. This recording was selected for good

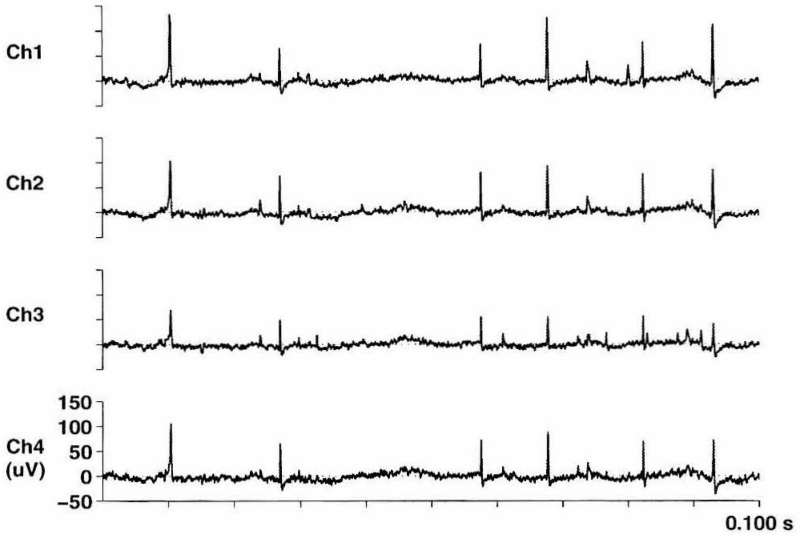

Figure 56: Sample Recording
An example recording made with the headstage in conjunction with the matching secondary amplifier of neural signals from a tetrode. Notice the action potentials (spikes) of varying heights across the four channels; each set of distinct heights is from a different cell. (cmem4714s.au)

isolation of a number of cells, but is by no means unusual or exceptional. Additional example data is given in Chapters 2, 3 and 6.

Plots of noise versus frequency for various input loadings can be seen in Figure 57. Resistors of 1 kΩ through 10 MΩ by decades were used in lieu of electrodes between the amplifier inputs and ground, adding the load resistance in parallel to the designed input filter at the first stage of amplification. The family of graphs depicts the noise spectrum for the range input loads: the set of curves all start with a common feature, namely a maximum at about 0.7 Hz, and depart from there with increasing frequency depending upon the input load. Data were recorded in a Faraday cage, with additional steps taken to minimize external interference, at a sampling rate of 25 kHz for 1000 s. Signals passed through a 10 kHz anti-alias filter before conversion, generating both the ripples at frequencies above 1 kHz, and the sharp cutoff at 10 kHz evident in the graph. Although not depicted for reasons of clarity, recordings were also made with input loads of 0 Ω (shorted inputs), and the spectra were indistinguishable from those with 1 kΩ loads.

If we assume the observed noise is solely due to thermal noise from the input network (including input load), we can predict the observed noise values using the following formula:

$$V_\eta = 1.29 \times 10^{-10}\sqrt{\Delta f \times R}, \tag{8}$$

where Δf is the bandwidth of the measured signal in Hz, R is the value of the resistor in Ohms, and V_η is the theoretical noise voltage expressed in Volts. The constant has been selected for

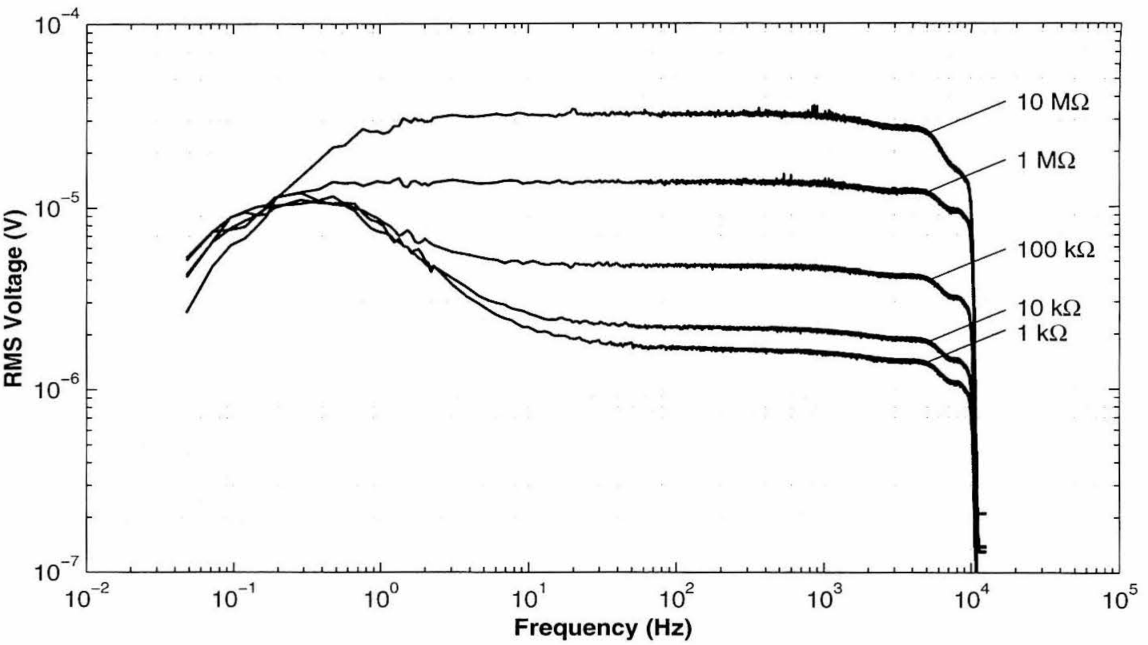

Figure 57: THA2 Noise Response
Input-referred noise versus frequency for input loads of 1 k, 10 k, 100 k, 1 M, and 10 MΩ.

R (Ohms)	$R \parallel 10^7$ (Ohms)	V_η (μV)
0	0	0.00
10^3	9.99×10^2	0.41
10^4	9.99×10^3	1.3
10^5	9.91×10^4	4.1
10^6	9.09×10^5	12.
10^7	5.00×10^6	29.

Table 5: THA2 Theoretical Noise Levels
This table shows the thermal, or Johnson, noise generated by resistors at room temperature. The leftmost column is the values of input loads, R, used to generate Figure 57. The middle column is the effective input load, the value of R in parallel with the value of the bias bleed resistors in the headstage input network (10 MΩ). The right column is the expected noise for the parallel resistance.

room-temperature operation, where expected values will be in the microvolt region. Assuming the bandwidth is 10 kHz (the cutoff of our anti-alias filter), the theoretical noise voltages, V_e is given for the various input loads in Table 5.

Comparing these noise values to the flat portion of the measured noise curves above 10 Hz, there is an excellent match for the highest input loads, and a progressively poorer match for smaller values. We conclude, therefore, that for input loads above approximately 100 kΩ, the observed noise

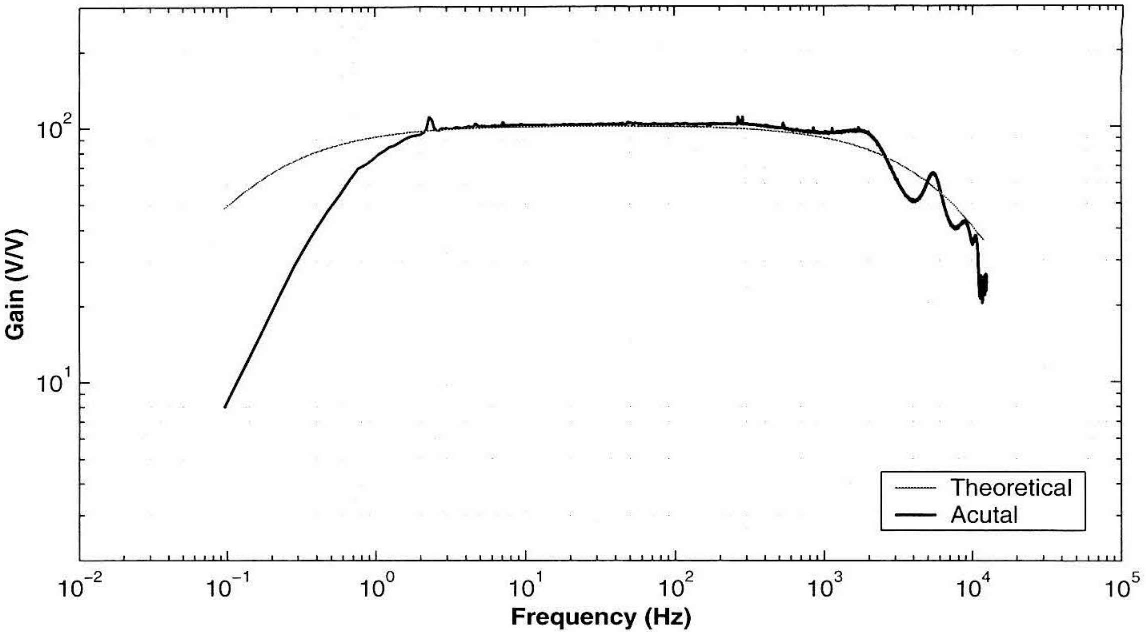

Figure 58: THA2 Gain Response
Gain versus frequency measured with white noise input of 0.1 mV RMS and an input loading of 1 MΩ.

is dominated by resistor thermal effects, and, below that value, other sources become primary. We might expect these other sources to be the current noise on the first opamp stage, or resistor noise from feedback elements in that and later stages. Further, the observed noise curves become more complicated below 10 Hz where additional mechanisms, including the frequency response of the input network, come into play. We might speculate that at lower input resistances, at frequencies above 10 Hz, we are observing effects from first stage feedback elements, which transitions below 10 Hz to $1/f$ input current noise and the effective larger input load due to input network filtering effects also from the first stage, and which, in turn, becomes suppressed below 1 Hz due to high-pass filtering in later stages of instrumentation. But this is just speculation, and these details will not be pursued further in the present document.

A plot of headstage gain over frequency can be seen in Figure 58. To generate this figure, a noise source of known amplitude was fed to the input of the headstage through a 1 MΩ resistor, and both the input and output were captured. The spectrum of the output was then normalized by the spectrum of the input, generating an estimate of the transfer function over frequency for the amplifier. The low-frequency response is compounded by later stages of amplification in the instrumentation equipment, and the ripples above 1 kHz are due to mismatching of the anti-alias filter responses. Effects from 60 Hz interference have been digitally removed with a non-linear filter. The theoretical trace on the plot is the designed response of the headstage.

Advanced Circuits
21100 E 33rd Drive
Aurora, CO 80011
800.289.1724
http://www.4pcb.com

Bi-Tronics
10 Skyline Drive
Hawthorne, NY 10532
800.666.0996
http://www.bitronics.com

National Semiconductor
http://www.national.com

Digi-Key
701 Brooks Avenue South
Thief River Falls, MN 56701
800.344.4539
http://www.digikey.com

Surface Mount Distribution
16321 Gothard Street, Unit G
Huntington Beach, CA 92647
714.841.4556

Linear Technology
http://www.linear.com

Table 6: THA2 Parts Sources
A listing of sources for the parts required to construct the THA2.

3.8 Improvements

Future versions of the headstage will use an inverting configuration for the second stage to improve high-frequency rejection, and provide an overall negative system gain when used with secondary amplifiers. The negative system gain is desirable to bring extracellular action potentials, normally negative spike-like excursions, into the positive realm for ease of visualization. Also, considerations are being made to include a fully differential signalling path.

3.9 Summary

A low-noise four-channel headstage amplifier for making neuroelectrical recordings has been presented. Technical specifications, measured performance, parts list, and sources have been supplied, allowing the reader to readily duplicate the design. For additional information, including construction advice, Gerber plots and the like, please contact the author. Commercial concerns please note that this design is part of a patent-pending system, and the author should be contacted regarding any for-profit uses.

3.10 Headstage Sources

A list of sources for the parts and supplies used in this design can be found in Table 6.

4 A Four Channel Secondary Amplifier

The second stage of amplification in an electrophysiology system is not nearly as critical as the first. Thus, the design constraints being less severe, this is often where filtering and variable gain are introduced. This section will present the design of such an amplifier module for use with Tucker-Davis Technologies equipment. This four-channel adjustable gain module is used in neural recordings to amplify signals prior to visualization and A/D conversion. The amplifier is low-noise, includes low and high pass filtering, and has settable gains from 1–1000 in factor steps of 1, 2, 5.

4.1 Introduction

The module described in this section is a four-channel amplifier for use in TDT-based systems. The design goals were to create a low-noise robust amplifier with gain adjustable from 1 to 1000 for all four channels in stepwise fashion from a single control. Additionally, the amplifier had to receive signals from an associated headstage, such as the THA2 described in a previous section, provide power to it and be insensitive to overloads both on signal and power lines. The input and output was to be buffered; the amplifier to provide high-pass filtering at 0.1 Hz, and low-pass filtering at 25 kHz. High-pass filtering is primarily to block DC; low-pass filtering is to suppress contamination from the eye position system which produces signals from 75 kHz on up. As limited power is available from the TDT system, power consumption was to be minimized, if possible. Finally, the gains were to be accurate to 1%.

While input signals were to be primarily neural-level as generated by the headstage (with its designed gain of 100), for the sake of flexibility and adaptability to future headstages, the full ±10 V range of input values common to TDT equipment was to be supported. Similarly, the output range was also to span the full ±10 V range.

4.2 Design Description

The schematics for the AMP4 module can be seen in figures 59–60. The amplifier is intended to have four low-noise channels with a ganged gain control, thus, the central feature of the design is a dual voltage-controlled amplifier (VCA), the Analog Devices AD600. Each channel has a Burr-Brown INA111 input buffer arranged with a 0.1 Hz high-pass filter and optionally grounded negative input to allow for single-ended or differential signalling. Two stages of adjustable amplification, allowing for potential gains in the range of 0–80 dB, are followed by an output buffer.

The choice of the AD600 was, in the end, a regrettable one, because of the significant design constraints it imposes. The amplifier runs off of ±5V power supplies meaning that the maximum signal range is well below the intended ±10V module output range. To accommodate the lower intermediate range, the signal is downconverted by a factor of five before delivery to the VCAs, amplified, and then upconverted by a factor of five (a factor of 5 was chosen to provide adequate headroom; a factor of 2.5 would likely have sufficed). The downconversion happens between the second and third stages, and the upconversion happens at the output stage. Additionally, the input networks on the AD600 have a specified resistance of 100 Ω, thus any device delivering signal to these chips sees a heavy load. Heavy loads translate into heavy power dissipation. The down/up-conversion translates into potential for premature clipping at high input signal levels. While the amplifier performs well in terms of noise and adjustability, it consumes a lot of power and the input

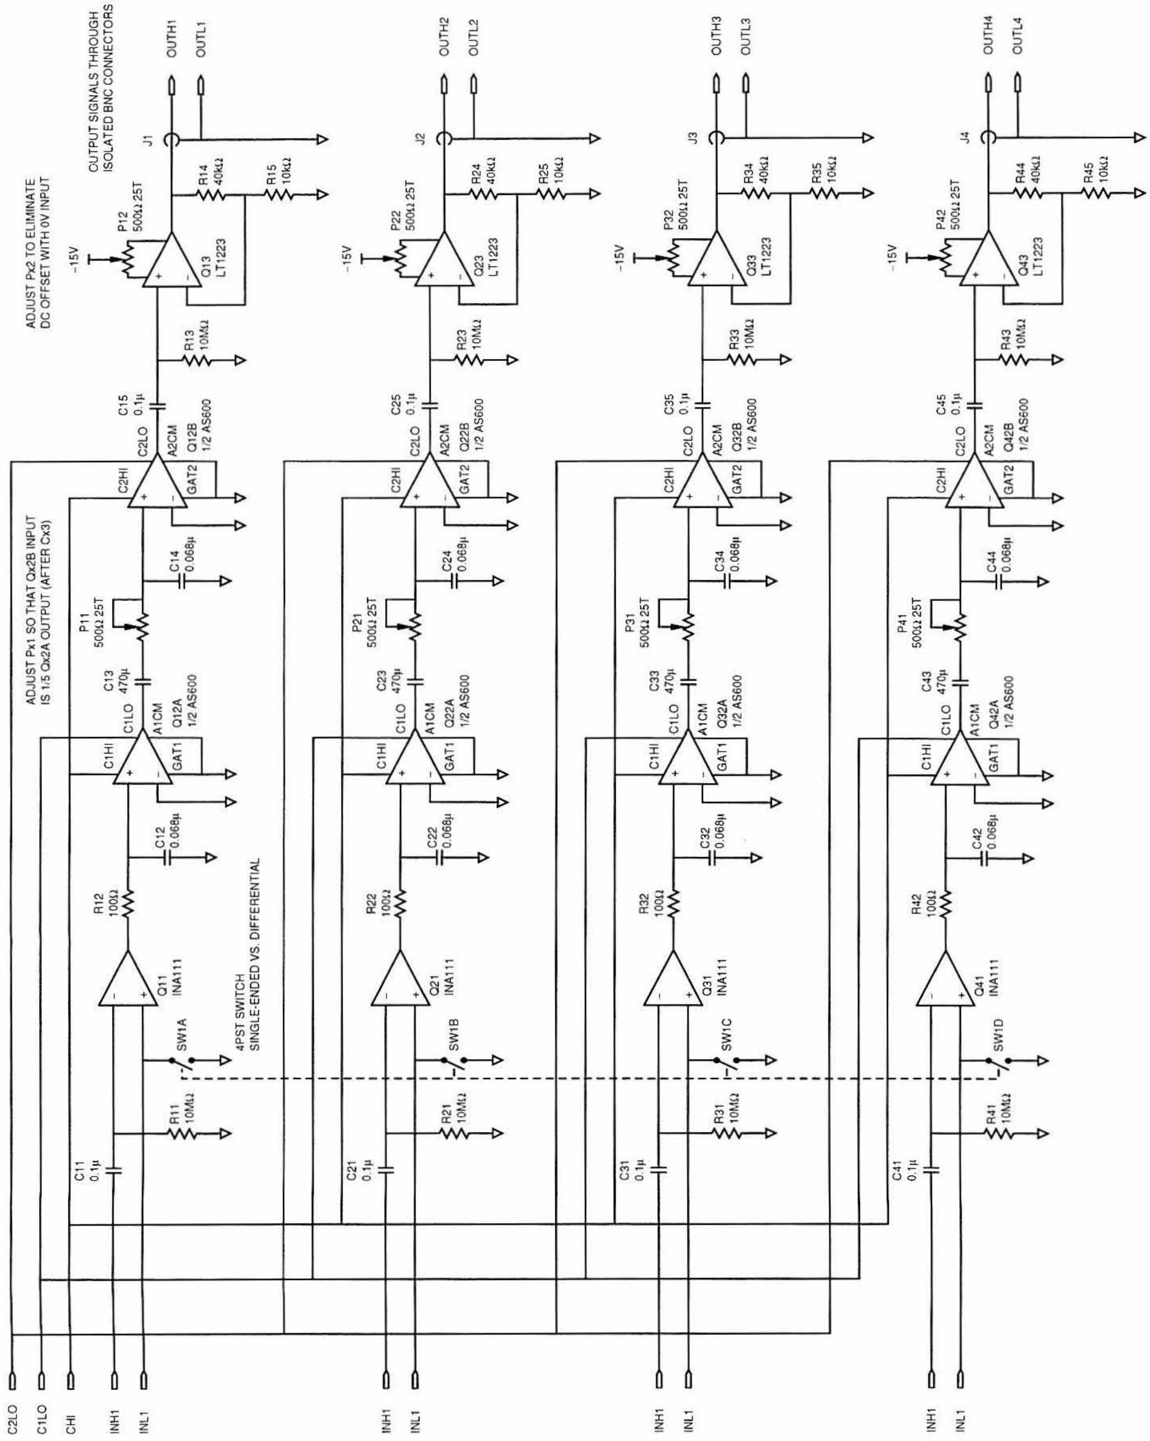

Figure 59: AMP4 Schematic 1: Signal Paths

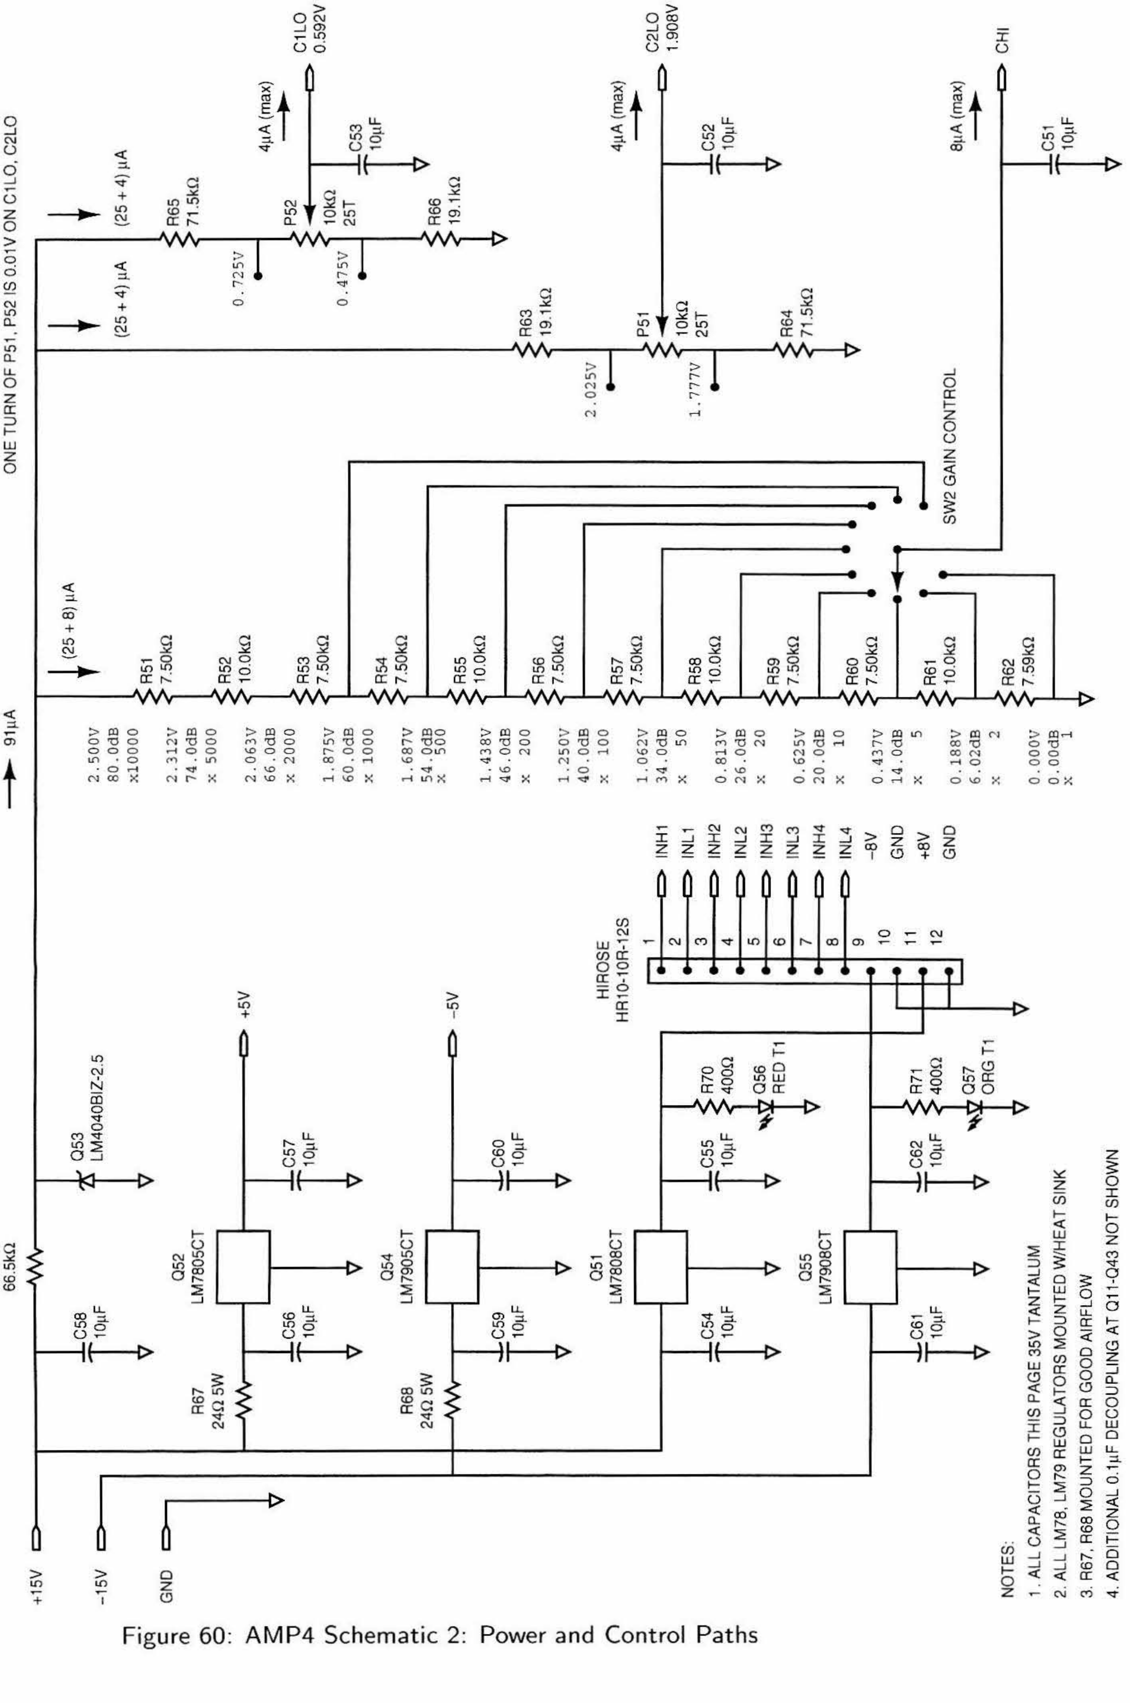

NOTES:
1. ALL CAPACITORS THIS PAGE 35V TANTALUM
2. ALL LM78, LM79 REGULATORS MOUNTED W/HEAT SINK
3. R67, R68 MOUNTED FOR GOOD AIRFLOW
4. ADDITIONAL 0.1µF DECOUPLING AT Q11-Q43 NOT SHOWN

Figure 60: AMP4 Schematic 2: Power and Control Paths

range varies with gain in a way that is counterintuitive.

In detail, the signal path is as follows. Four channels are arranged identically in four op-amp stages each. These four stages are conceptually arranged as three: an input buffer, a two-stage voltage-controlled amplifier, and an output buffer. These will be discussed in turn from the input to the output.

The first stage receives the signal from the headstage, buffering and filtering it. The input is received by an RC filter with $f_c = 0.1$ Hz that also provides a bleed path via a 10 MΩ resistor for the bias current. The stage provides no gain, but can be switched between single-ended and differential mode. The intent is that the signal is generated at the headstage in single-ended mode, but received at the secondary amplifier in differential mode; this is often called pseudo-differential, and provides increased common-mode rejection over single-ended systems while not having the additional component count of full differential systems. The output of the first stage drives a 100 Ω current limiting resistor in series with the input to the second op-amp, part of the voltage-controlled amplifier. The current limiting resistor is intended to reduce the severity of the load presented to the output of the first stage.

The second stage is formed by two halves of an AD600 which together form a cascaded voltage-controlled amplifier. The gain across the two op-amps is controlled by the voltages on three lines: two lines, C1LO and C2LO form offsets for the first and second stages, respectively, and the third, CHI specifies the combined gain. For increasing voltages on CHI, the two stages are activated in sequence: the first stage increases up to a maximum gain of 40 dB before the second stage is engaged; once the first stage is at maximum gain, for continuing increases in CHI, the second increases to its maximum of 40 dB. Filtering is placed between the two stages, a 0.1 Hz high pass filter, a voltage divider of 1/5, and an additional 25 kHz low pass filter. The fixed reduction in gain from the voltage divider is compensated by the gain in the final stage, and also reduces the loading on the output of the first half of the AD600.

The third stage has a final 0.1 Hz low-pass filter, adjustment for DC offset, and a gain of five. The gain at this final stage insures that the overall gain can be set to the desired values as indicated on the front panel. Note, however, the apparent uncompensated factor of 1/2 lost between the first and second stages, and compare against the measured gains in Table 7.

As shown in Figure 60, the control voltages for CHI are generated by precision voltage divider from a 2.5V reference, and selected through a front panel rotary switch. A capacitor on the control line insures stability and reduces the influence of external noise coupled to the line. The two offset voltages C1LO and C2LO are generated in a similar fashion through multi-turn potentiometers. Current on the control lines and dividers is designed to remain well within the limits of the voltage reference.

Also shown in Figure 60, power is locally regulated from the provided ±15 V down to ±5 V

item	*manufacturer; part number*	*description*
PC1	custom	printed circuit board
Q[1–4]1[†]	Burr-Brown INA111	instrumentation amplifier
Q[1–4]2	Analog Devices AD600	voltage-controlled amplifier
Q[1–4]3	Linear Tech LT1355	high-speed op-amp
Q51	LM7808CT (TO-220)	8V positive regulator
Q52	LM7805CT (TO-220)	5V positive regulator
Q53	National Semi LM4040BIZ-2.5	2.5V reference
Q54	LM7905CT (TO-220)	5V negative regulator
Q55	LM7908CT (TO-220)	8V negative regulator
Q56		T1 red LED (+8V indicator)
Q57		T1 orange LED (−8V indicator)
R[1–4]1,R[1–4]3		10MΩ resistor
R[1–4]2		100Ω 1% resistor
R[1–4]4		40kΩ 1% resistor
R[1–4]5		10kΩ 1% resistor
R51,53,54,56, 57,59,60,62		7.50kΩ 1% resistor
R50,52,55,58,61		10.0kΩ 1% resistor
R63,R65		19.1kΩ 1% resistor
R64,R66		71.5kΩ 1% resistor
R67,R68		250Ω 5W resistor
R69		66.5kΩ 1% resistor
R70,R71		400Ω resistor
C[1–4]1		0.1μF capacitor
C[1–4]2		0.068μF capacitor
C[1–4]3		470μF capacitor
C[1–4]4		0.068μF capacitor
C[1–4]5		0.1μF capacitor
C[1–4]6–12		0.1μF decoupling capacitor
C54–C62		10μF 35V tantalum capacitor
J1	Hirose HR10-10R-12S	12 pin connector
J[1–4]1		isolated BNC bulkhead jack
P[1–4]1,P[1–4]2		500Ω 25 turn potentiometer
SW1		4PST switch
SW2	Grayhill 71BDF30-01-1-AJN	10 position rotary switch
H1–3		TO-220 heatsink

Table 7: AMP4 Parts List
A list of parts for constructing one amplifier. When the manufacturer, part number, or both are missing, it indicates the exact source is unimportant. All non-tantalum capacitors are ceramic chip NPO or X7R series. [†] *The notation, eg, Q[1–4]1 signifies the set of parts Q11, Q21, Q31, and Q41.*

for the AD600 chips as well as to ±8 V to send to the headstage. Since the power supply current required by the voltage controlled amplifiers is quite high, 24Ω dropping resistors are placed before the ±5 V regulators to absorb some of the power dissipation due to the 10 V drops.

4.3 Parts List

A list of required parts is given in Table 7. At the time of this writing, these parts are all readily available, and reasonably inexpensive, with the exception of the AD600 which runs approximately $30 per chip. Some parts have not been included on the schematic, but are included in the table for completeness.

4.4 Operation and Performance

As can be seen in the photographs in Figure 61, the number of controls and connectors is small, and therefore operation is straightforward. Connections are made on the front panel to a compatible headstage via the input connectors and a recording device or oscilloscope through output connectors. Gains are set through the front panel rotary switch, and receiving mode (single-ended or differential) is selected. A slight delay accompanies the selection of a new gain as the RC time constant of the control voltage CHI settles to a new value, and the high-pass filters adjust to new DC blocking levels.

A plot of measured gains for a set of four constructed amplifiers is shown in Figure 62, and Table 8. The gains are, broadly speaking, under the designed values by approximately a factor of 2; this suspiciously matches the uncompensated voltage divider between the first and second stages of amplification. Even considering this factor, the gains are, however, quite far from the intended 1% accuracy; the source of the observed errors has not been identified.

A plot of typical gain versus frequency is shown in Figure 63 for a commanded gain of 100. The roll-offs from high-pass and low-pass filtering are clearly visible: the measured high-pass corner frequency at 3 dB down is $f_c = 0.4$ Hz, significantly above the designed 0.1 Hz, but the multiple zeros broaden the corner such that the nominal 3 dB point does not adequately characterize the roll-off. The low-pass 3 dB point is right at the designed 25 kHz, leading one to believe that a mistake was made during design regarding the low frequency filter. Indeed, careful examination of the three high-pass filters reveals cutoff frequencies of 0.1 Hz, 0.85 Hz, and 0.1 Hz, respectively, a mixture of which should generate the observed f_c.

The expected filtering effects can be expressed as:

$$h = h_1 h_2 h_3 h_4 \tag{9}$$

$$h_1 = \frac{R_1 C_1 s}{R_1 C_1 s + 1} \tag{10}$$

$$h_2 = \frac{R_i}{R_i + R_i(R_2 C_2 s + 1)} \tag{11}$$

$$h_3 = \frac{R_i C_4 s}{R_i C_3 s + (R_p C_3 s + 1)(R_i C_4 s + 1)} \tag{12}$$

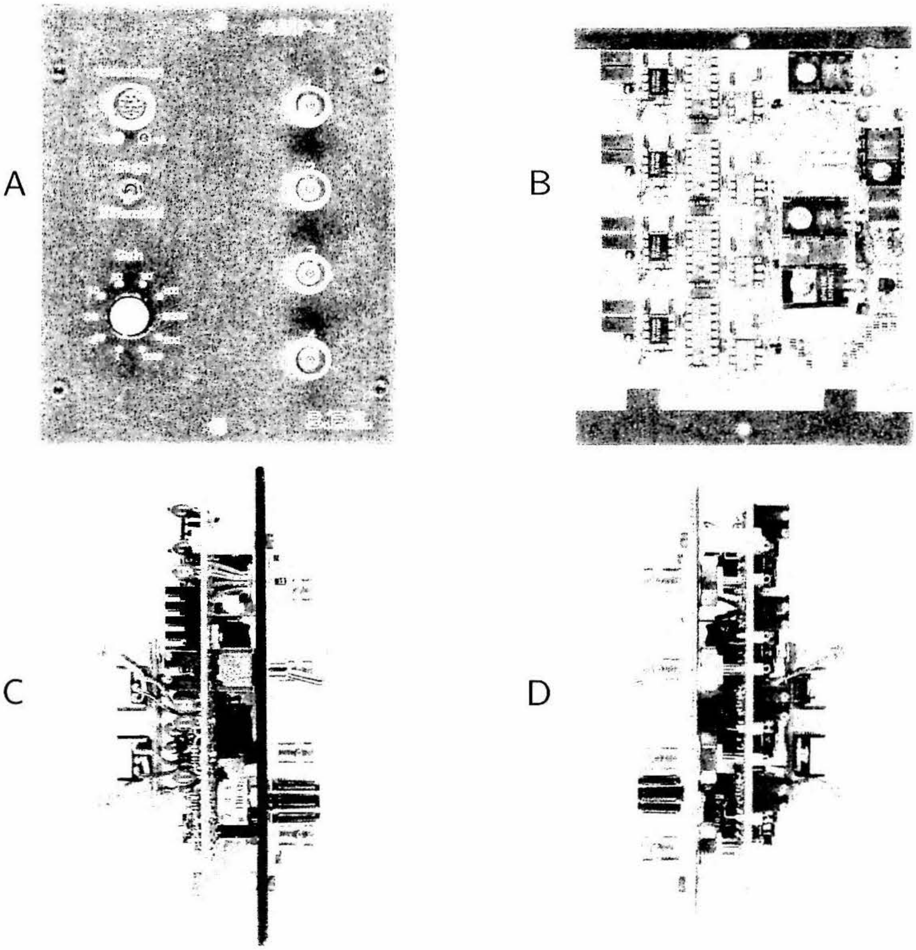

Figure 61: Four Views of Assembled AMP4 Module
A, *front view. The headstage connector is at the upper left, the receiving mode switch in the middle left, and the gain control knob on the lower left. The four BNC output connectors are on the right.* **B**, *rear view. Signals are flowing right-to-left in this photograph. The power regulators can be seen on the right, with their associated heatsinks, and the voltage divider resistors for the control voltages. Down the middle are the INA111 receivers; directly to their left are the AD600s, and to their left, the output opamps. At the far left the 25 turn potentiometers can be seen arranged pairwise.* **C**, *left side view. The small daughter card built to hold the* ± 5 *V regulators and* 24 Ω *dropping resistors can be seen.* **D**, *right side view. Another view showing the daughter card and output connectors.*

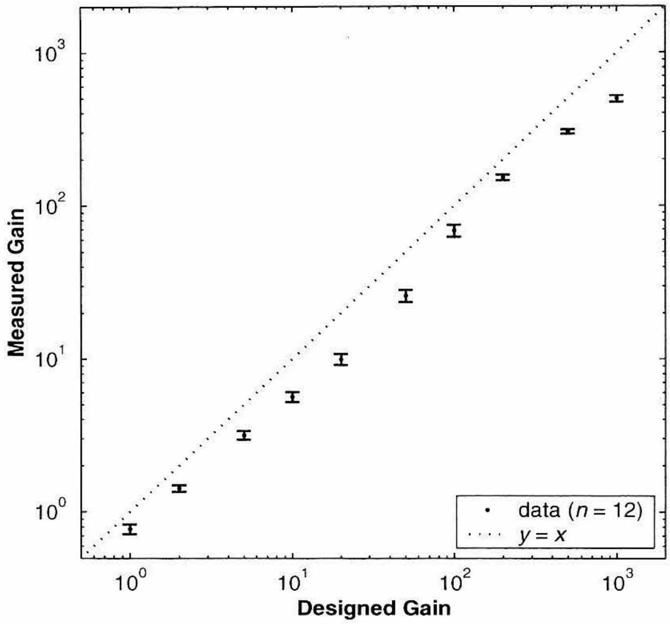

Figure 62: AMP4 Amplifier Gains
Gains are shown as measured by examining a broadband noise signal before and after amplification. Signal was generated as uniform noise at 1.5 V peak-to-peak, attenuated by 40 dB, and fed to the amplifier through 1 kΩ resistors. The unattenuated and amplified signals were then anti-alias filtered (10 kHz), digitized at 50 ksamp/sec, Fourier transformed, averaged over 10–1000 Hz, and the ratio (amplified to normal) taken to compute the gain. Data points represent channels 2–4 of from each of four amplifiers (the channel 1s were discarded due to the necessity of recording the unamplified input signal).

designed	*measured*
1.00	0.77 ± 0.057
2.00	1.4 ± 0.072
5.00	3.2 ± 0.20
10.0	5.6 ± 0.42
20.0	9.9 ± 0.82
50.0	26. ± 2.3
100.	69. ± 6.2
200.	151. ± 6.2
500.	301. ± 9.6
1000.	495. ± 25.

Table 8: AMP4 Amplifier Gains
These are the values plotted in Figure 62; see that caption for collection details. The mean error factor is roughly 2/3. For $g = 200$, the most commonly used gain, the error factor is slightly closer to 1, but well off the intended 1% error.

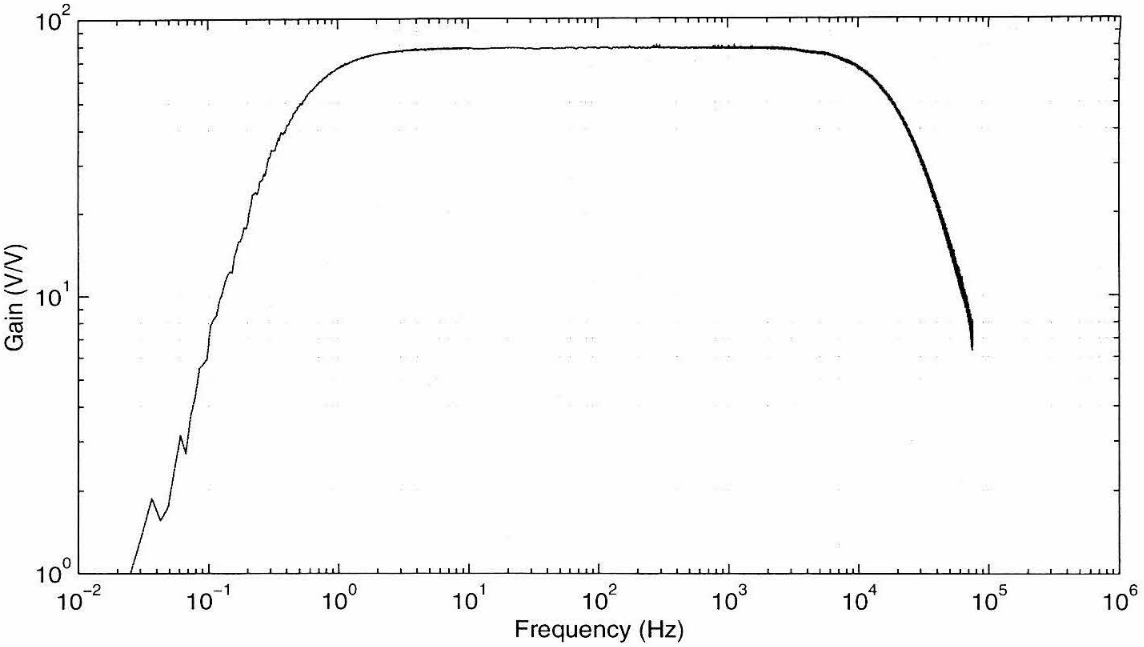

Figure 63: AMP4 Gain versus Frequency
This is a typical frequency response curve for one channel from an amplifier at a setting of $g = 100$ showing gain on the vertical axis and frequency on the horizontal axis.

$$h_4 = \frac{R_5 C_5 s}{R_5 C_5 s + 1} \tag{13}$$

where, for example, R_1 can be any of R11, R21, R31, R41, and R_i is the input resistance of the AD600. When combined, and normalized for system gain, these produce the reference trace on Figure 63.

While this design has serious shortcomings in light of the original goals and intended performance, it does meet many of the goals, including reliability and robustness. For example, although not shown, the response to overload on headstage power draw (for example, with a headstage failure) is shutdown on the ±8 V regulators without affecting the rest of the amplifier or TDT system. The response to overload on the input signals similarly goes to a saturated state (which is DC-decoupled to 0 V), followed by a graceful recovery once the overload is removed.

However, the shortcomings have been addressed with an improved version of the amplifier which has been designed but not implemented. Should the lowest frequencies, for example, below 1 Hz, become experimentally important, it will become imperative to construct the improved version.

4.5 Summary

In brief form, we have covered a module that was constructed to amplify low-level signals. The design here misses the goals of reasonable power dissipation, full signal range, and 1% gain accuracy, but it achieves the goals of reliability, insensitivity to overloads from the headstage, and correct operation for neural-range signals.

5 A Double-Throw Four-Pole Electrically Controlled Switch

While not a critical part of the signal path in terms of conditioning, it is important to have a means for synchronizing recorded neural signals with behavioral events. This section will present the design and construction of a switching module for use with Tucker-Davis Technologies equipment. This quad single-pole-double-throw (4×SPDT) module is used, for example, to periodically switch between a recorded voltage and a reference voltage during neural recordings to mark the time of external events.

5.1 Schematic and Design Description

The schematic for the switch module can be seen in Figure 64. The straightforward circuit has four independent single-pole double-throw switches where each switch is controlled by a logic-level voltage. Inputs and outputs are labeled in banks, signal inputs A_i, B_i, control inputs C_i, and outputs D_i, where i ranges from 1 to 4. The behavior of the circuit is described by the following equation:

$$D_i = \begin{cases} A_i, & \text{if } C_i = 0 \\ B_i, & \text{if } C_i = 1 \end{cases} \tag{14}$$

The analog switch chosen, the Maxim MAX333A, is fast, robust, and has low on-resistance. If the MAX333A is not available, the older MAX333 is a good alternate, as is the MAX4533. They are all pin-compatible. Switching times are well under 200 ns, and power consumption is well under 10 mW.

5.2 Parts List

A list of required parts is given in Table 9. At the time of this writing, these parts are all readily available, and reasonably inexpensive. The table lists IC1 and IC2 as the MAX333 sourced from Digi-Key, but the MAX333A is available directly from Maxim in sample quantities.

5.3 Construction

Most of the assembly time for the switch will be in wiring the circuitry, but a fair fraction will also be taken by drilling the necessary holes (some 24 in the front panel, and 16 in the circuit board). A drawing is provided in Figure 65 to assist in this, as well as photographs of the finished product in Figure 66. The BNC connectors should be tightened with sufficient torque to insure that they do not loosen over time, as tightening them after the module has been assembled is difficult. Care should be taken to route wires cleanly. Although the schematic as presented does not include power supply decoupling capacitors for the two integrated circuits, they should be included, one per voltage rail,

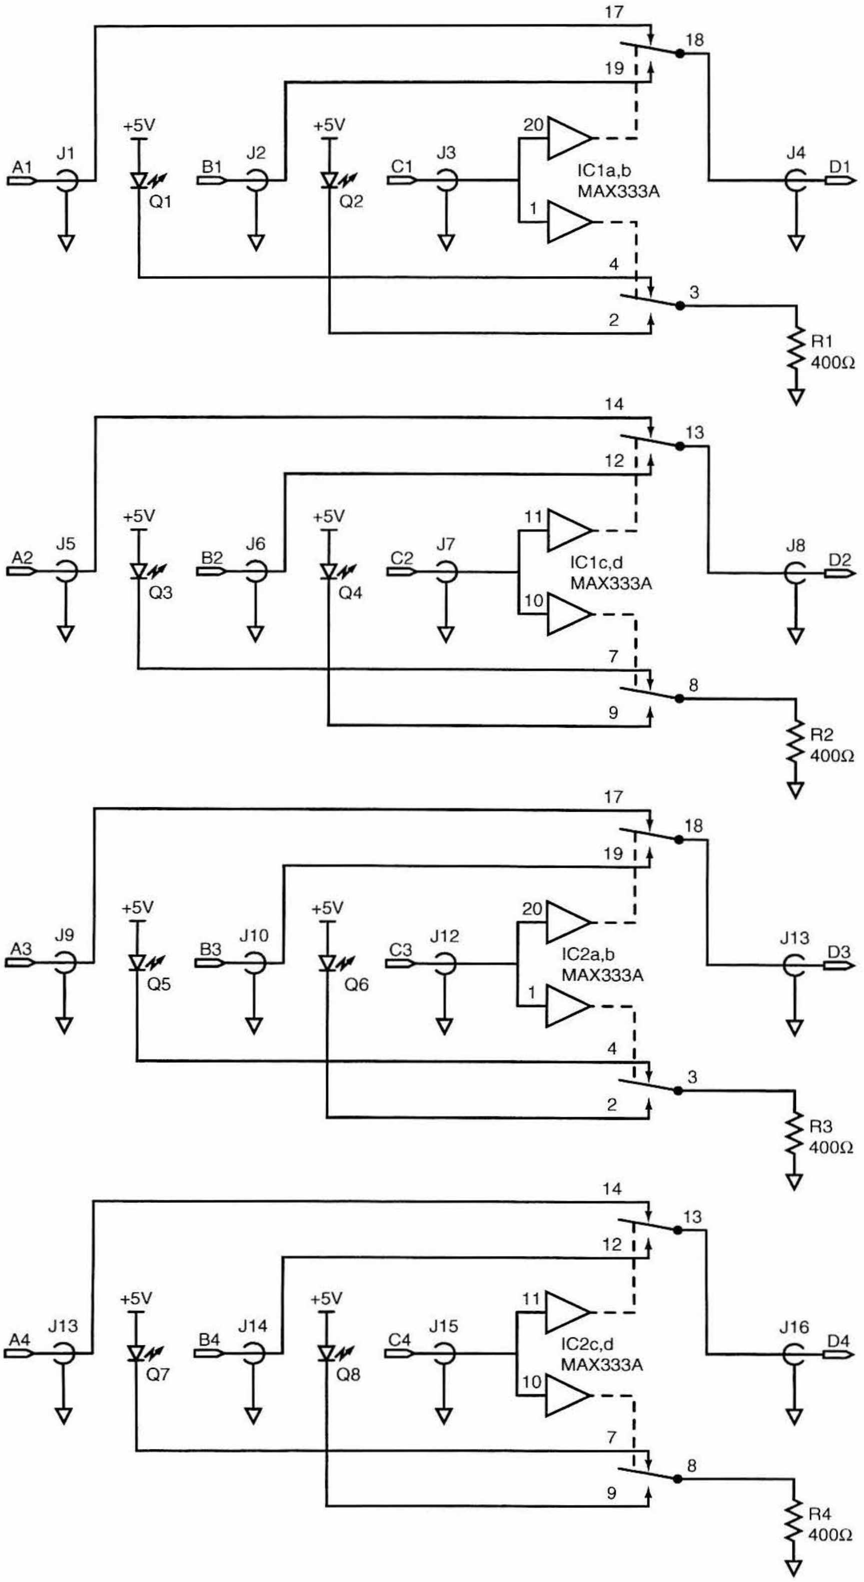

Figure 64: SW1 Quad SPDT Switch Schematic

item	*manufacturer; part number*	*description*	*source; order number*
BB1	TDT BB1	breadboard module	TDT BB1
IC1–IC2	Maxim MAX333CPP	quad analog switch	Digi-Key MAX333CPP
Q1–Q8	Chicago Miniature CMC01G	T1 green LED	Digi-Key CMC01G
R1–R4	Yaego MFR-25FBF 402R	400Ω resistor	Digi-Key 402XBK
J1–J16	Amphenol 31-10-RFX	isolated BNC connector	Digi-Key ARFX1063
C1–C4	Panasonic ECU-S1H104KBB	0.1μF 50V X7R	Digi-Key P4923
misc	Vector R32	press-in pins	Digi-Key V1059

Table 9: SW1 Parts List
A list of parts for constructing one switch.

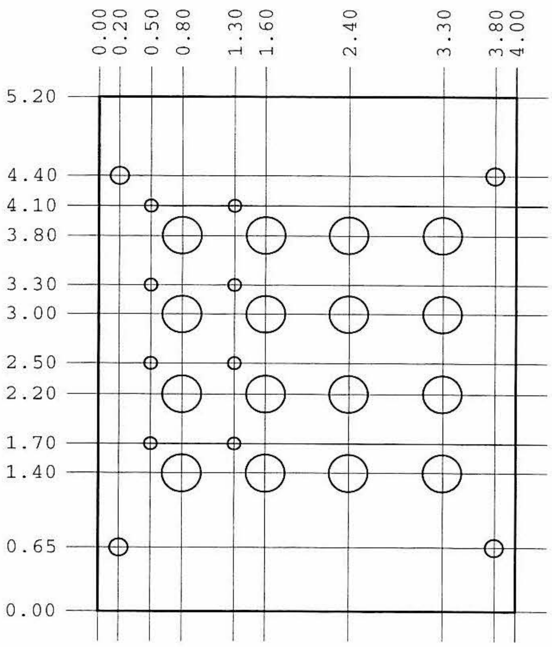

Figure 65: SW1 Front Panel Drawing
A layout of the holes to be drilled in the front panel. The four medium-sized holes are pre-drilled in the BB1 front panel for mounting screws. The eight small holes (0.125 inch) are for rear-mounted T1-size LED lamps; when assembling the board leave just enough lead length so that the LEDs are pressed up slightly against the front panel. The 16 larger holes (0.375 inch) are for the BNC connectors, and, if possible, should be punched as D-shaped cutouts.

adjacent to the packages. Additionally, there should be a larger set of capacitors at the module power entry point.

5.4 Operation

Operation is numbingly simple. Inputs are connected to A_i and B_i, control voltages to C_i, and outputs taken from D_i. The lights next to each A_i and B_i input are illuminated when the corre-

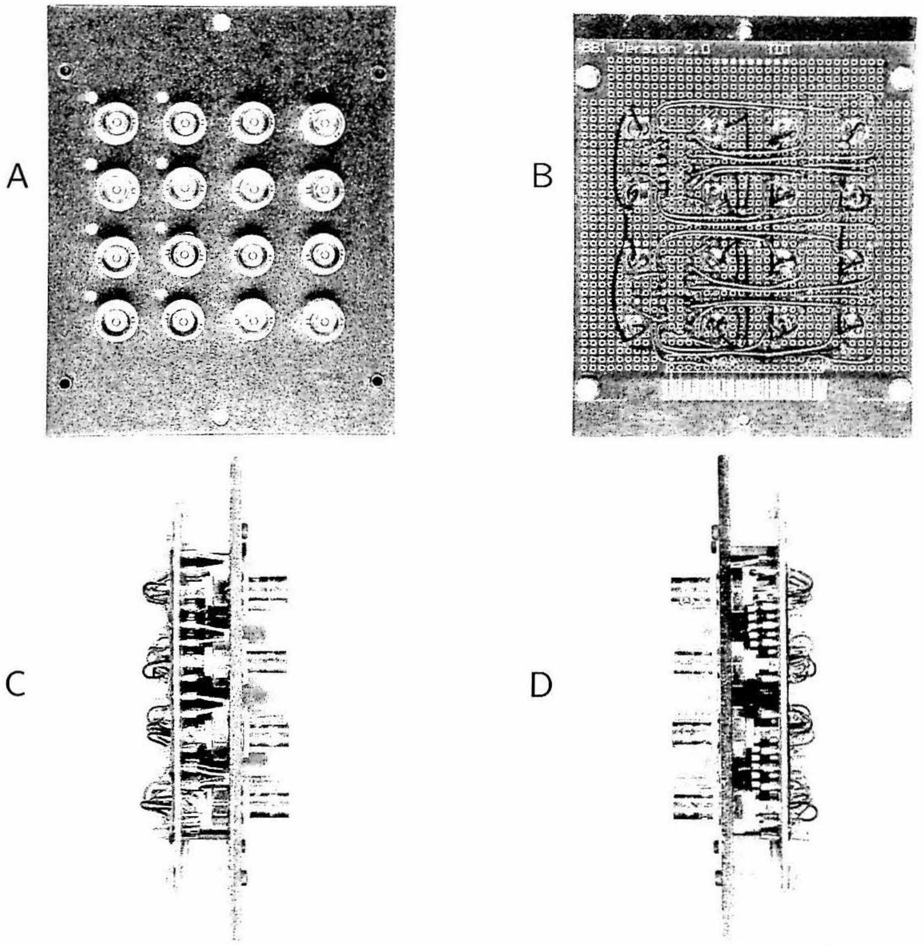

Figure 66: Four Views of Assembled SW1 Module
A, *front view. Going left to right the four columns are the connections for* A_i, B_i, C_i, *and* D_i, *respectively. The indicator LEDs can be seen above and to the left of each* A_i *and* B_i *input.* **B**, *rear view, showing wiring. Notice the holes made in the circuit board to allow access to the BNC connector soldering leads.* **C**, *left side view. The LED legs for the* A_i *inputs can be seen from this side.* **D**, *right side view. The MAX333 chips can be seen just behind the column of output BNC connectors.*

sponding input is being routed to the output. Signals A_i and B_i can range from −10 to 10 V, and controls C_i should be TTL/CMOS logic levels between 0 and 5 V.

5.5 Results

The results are nearly as simple as the operation. The switch is quiet, quick, and reliable. In use, we typically route signals from a multi-channel neural recording to the A_i inputs, put a grounding cap over the B_i inputs, and have a computer-driven logic level feeding all four C_i inputs in parallel. This is used to periodically impress a fixed voltage (0 V for 10 ms) on the recording by the behavioral control computer as can be seen in Figure 67; these marks can be used to later synchronize the

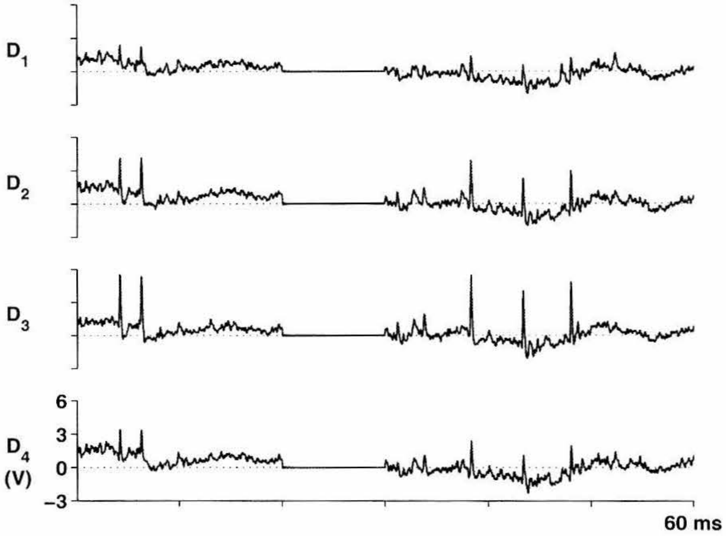

Figure 67: Sample Switching
An example recording switching between four channels of neural signal on A_i and ground on B_i. The vertical scale is signal strength in volts for each of the four outputs from the switch; the horizontal axis is time in milliseconds. D_i was switched from A_i to B_i at 20 ms, and back 10 ms later. (cmem3310s.au)

neural signal to other records. Although not presented here, these values have been readily and repeatedly recovered by an inverse threshold (signal must be in magnitude below a certain value determined by the inherent noise in the A/D subsystem). Precise characterization of the switch is beyond the scope of this brief paper.

5.6 Summary

In brief form, we have covered a simple module that was constructed to allow remote electronic selection between two banks of signals. The parts are readily available, construction is straightforward, and the results have been excellent.

6 Chapter Summary

This chapter reviewed the instrumentation used to capture data in this dissertation, providing sufficient detail on three of the constructed subsystems, the tetrode headstage amplifier, the secondary amplifier, and a switching module, for interested researchers to construct their own copies. These systems were all used locally to collect neural data, streaming the voltages in continuous fashion to archival storage. Conversion from this continuous stream into separated spike trains is the subject of the next chapter.

Although dc amplifiers can provide accuracy, bandwidth, and gain, it is not as simple as just purchasing instrumentation equipment and connecting it between transducers and recorders. It is, unfortunately, more complicated.

— RALPH MORRISON (Grounding and Shielding Techniques in Instrumentation, 1967)

Physiological processes can be a bear to measure and evaluate.

— LAURA HELMUTH (NetWatch, *Science*, 31 March 2000)

Chapter Six: Spike Sorting

As a main focus of his doctoral work at Caltech, and as part of our collaboration, Maneesh Sahani developed an advanced spike sorting algorithm [67]. While broadly applicable, it was specifically aimed at sorting single unit signals from continuous voltage tetrode recordings. This chapter will discuss the implementation and application of the algorithm, and while we were both involved in this effort, the text herein should by rights bear Maneesh's name as first author.

1 Introduction

The question this chapter aims to answer is how we were certain of the isolations in our recordings given the claim made of 2–5 cells per site. Towards this, we will discuss the mechanisms used. This will start with a description of the data structures, followed by a review of the algorithm, and some example results.

The spike sorter was written using MATLAB, which provides an excellent development environment. While the resulting code is perhaps not as efficient as it possibly could be, for example, many of the computations could be performed on 16-bit integer data rather than double-precision floating point, it is portable and readily modified to incorporate new features or correct problems with old ones. Optimizations have been made where possible, and on 450 MHz Pentium II processors, contemporary computers when the project was started, it runs reasonably quickly.

2 Data Structures

The data structures are a design that evolved out of conversations with Carlos Brody, and, later, work with Jennifer Linden. One `expt` data structure is created for each block of trials, conceptually an independent experiment. The idea is to create a framework for recording experimental data which allows efficient data manipulation for analysis, while also being adaptable with little additional work to many paradigms. The topmost fields within the structure are the experiment name, the date performed, an array of structures, one for each trial of the experiment, a status vector indicating success or failure for each trial, and a cell array of validly sorted units (cells) for the experiment.

Each element of the trial array is a structure which contains the start and end times, an array of spike times, cell models, and behavioral markers. The behavioral markers are different for each paradigm, but are expected to be identical for each trial. If there are multiple trial conditions with varying numbers of markers, then they would need to be gathered either in a matrix, a structure, or a cell array.

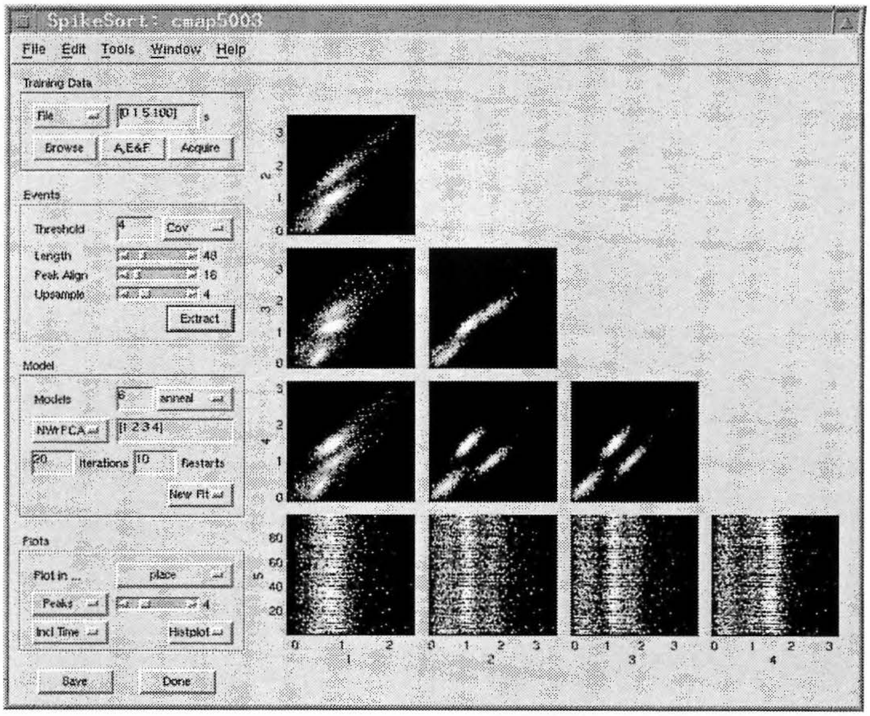

File [0 1 5 100] s

Threshold 4 Cov

Length 48

Peak Align 16

Upsample 4

Models 6 anneal

NWrPCA [1 2 3 4]

20 Iterations 10 Restarts

Plot in ... place

Peaks 4

Incl Time

Figure 68: GUI for Spike Sorting

This is a screen shot of an example recording (cmap5003) being processed. When taken, the data had been loaded in 100 segments of 1 second each, spaced every 5 seconds, and events extracted. The event threshold was set to 4 times the covariance, and 48-point (by 4 channels) events formed with the peak aligned to the 16th position after resampling to 4 times the base sampling rate. The dimensions plotted in the clustergram are the peak voltages across the four tetrode tips and time of occurrence.

3 Overview of Algorithm

The algorithm has been implemented in two stages. In the first stage, models of the detectable spikes in a recording are built using a relatively short segment, and in the second stage, filters developed from these models are applied to the entire recording. The reason for this architecture is to allow for on-line processing where a brief, perhaps two minute, recording would be made to train the models, after which an incremental version, even if not truly on-line, would parse the recording trial by trial, retraining as necessary.

3.1 Phase One

The generation of spike models has three steps to it. These steps are most easily performed through the graphical user interface called `SpikeSort`. The three steps are loading data from a file, extracting events from the data, and clustering the events into models. A screen shot of a recording being processed is shown in Figure 68, after the first two steps have been completed.

The first step loads the data from a recording into main memory and high-pass filters it. The

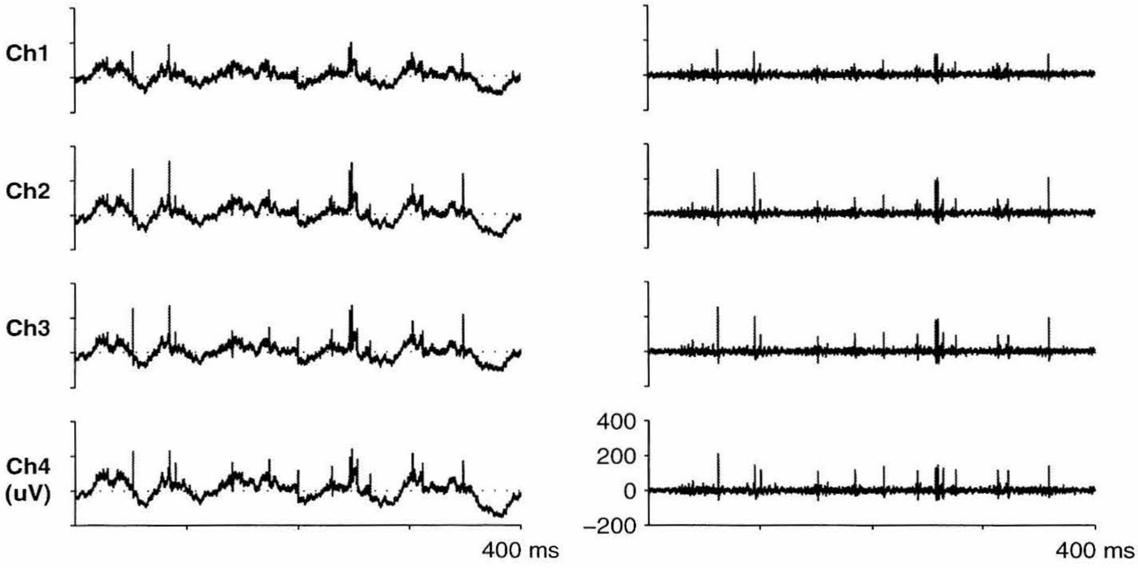

Figure 69: Data Before and After Filtering
The left panel shows a 400 ms segment from the example recording used in this chapter, as captured. The right panel shows the same 400 ms segment after digital high-pass filtering, as is used during event extraction and sorting. Vertical scales on the two panels are identical.

filtering parameters are stored in a structure called `SS_parameters`, and can be modified by hand; nominal values are $f_c = 600$ Hz and $n = 256$ points. For firing rates in the range of 10–50 Hz, such as found in LIP, a segment up to 2 minutes of total time is loaded. For off-line analysis, this is often specified as a sampling of 1 second segments spread out over the entire recording. Spreading the segments allows compensation for drift or other slow temporal variations in spike shape.

The second step extracts events from the loaded data. The loaded data are whitened in time and in space (across channels), or *sphered*. Then, excursions which are 3–5 times the covariance for a minimum of 2 samples are detected and used to extract a short segment of the loaded recording. This is similar to the operation of many commercial spike sorters, with the exception of the sphering process. Then, each extracted segment is upsampled by Fourier reconstruction typically by a factor of 4 before the center of mass of the event that lies above 2 times the covariance is aligned to a reference position, normally 1/4 or 3/8 of the way through the segment. Following this, the segment is decimated back to the original sampling rate and saved for further processing. The alignment step is important because it helps reduce a sampling artifact that creeps in from our relatively low sampling rate wherein the true peak might not have fallen on a sample point. The noise thus introduced is uniformly distributed making it difficult to model. For a fuller explanation, please see Sahani's thesis [67].

The third step applies a variant of the Expectation-Maximization (EM) algorithm called REM-

2 to discover the latent variables in the recording, which we hope to be cell identity. This clustering algorithm combines deterministic annealing (related to the more popularly known simulated annealing which in contrast is a stochastic algorithm) with EM to identify cluster centers. Normally, we apply the algorithm to the extracted data segments after they have been sphered. This has the effect of turning any variations that are due to background Gaussian processes into spheres, while leaving non-Gaussian variations due to bursting or drift non-spherical, greatly easing identification of the latter. Also, a rotational transformation is applied to the data which maximizes remaining variance and creates an approximation to the optimal linear discriminant. These seemingly arcane steps not only support the automated identification of clusters by mechanical means, but allow the accurate classification of events which would otherwise be considered noise.

REM-2 uses Gaussian clusters to model events, a background uniform cluster to mop-up any outliers, and a noise cluster at the origin to take responsibility for low-amplitude events which are likely to be unsortable. The algorithm starts with relatively large cluster sizes and iteratively shrinks them. After each step, EM is run to reassign events and the distribution of events assigned to each cluster is evaluated for Gaussianness with a Kolmogorov-Smirnov test. Each cluster which fails this test has its model cloned, and the events randomly partitioned between the original model and the clone. In simplified form, termination occurs when the cluster sizes reach the background noise (the actual termination condition is substantially more complex, and interested readers continue to be referred to Sahani's thesis).

At the end of this process, the generated models are placed in a structure called `SS_model`. This is then used in sorting the full recording.

3.2 Phase Two

The application of spike cluster filters to the recorded data stream happens on a per-trial basis. A file containing trial start times is required, and for the data and experiments presented in this dissertation would be generated by applying the tool `authreshold` to the recording. Each trial's segment of the recording is loaded from a file, filtered, and events extracted using the same mechanisms and parameters as in the model creation phase. If retraining has been enabled, the noise floor is remeasured and models are adjusted. Then, each event is projected into the model space and posterior probabilities generated against each cluster. The results are saved in the per-trial part of the experiment structure, including a vector of all spike times, a vector of highest probability assignment, a matrix of the full posteriors, and for each model, the times of spikes assigned to that model and the posterior probabilities for each. The spike waveforms are, at this point, discarded for memory storage considerations, although the structure could be augmented to include them. This part of the second phase is performed by the function `ss_sort` which returns an `expt` structure.

After the spikes have been extracted, the behavioral markers are merged into the result. The function for this will vary with experimental paradigm, but for this dissertation, the two merging

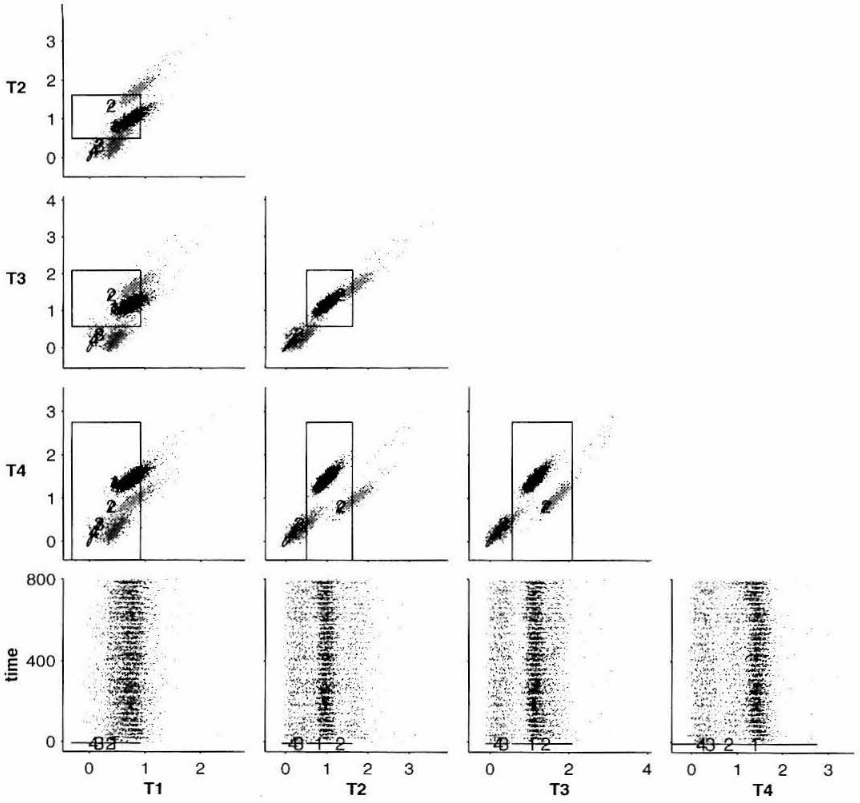

Figure 70: Sorted Events, Peak Space

functions are called **`hydload_memsac`** and **`hydload_maprf`**. These functions are based on a suite of MATLAB functions which read the binary data files created by our behavioral control program, Hydra, written by Maneesh Sahani.

4 Clustergrams

The results of a spike sorting are best viewed as clustergrams. Clustergrams, here, are series of two dimensional projections of higher dimensional spaces where each point in the space is plotted either in density form, as seen in the graphical interface figure, or scattergram colored with an assignment from a sorting. The triangular set of subplots is from the normally four dimensions of visualization, either the peak values on each tetrode tip, principle components, or noise-whitened principle components, plotted against each other. The bottom row, when present, shows the four against event time. A sorted clustergram is shown in Figure 70 for peak values, and in Figure 71 for noise-whitened principle component values.

Each point in the scattergram cluster plots is color-coded according to assignment, and each model is indicated by a numbered ellipse. In our example, the clustering was performed in the noise-

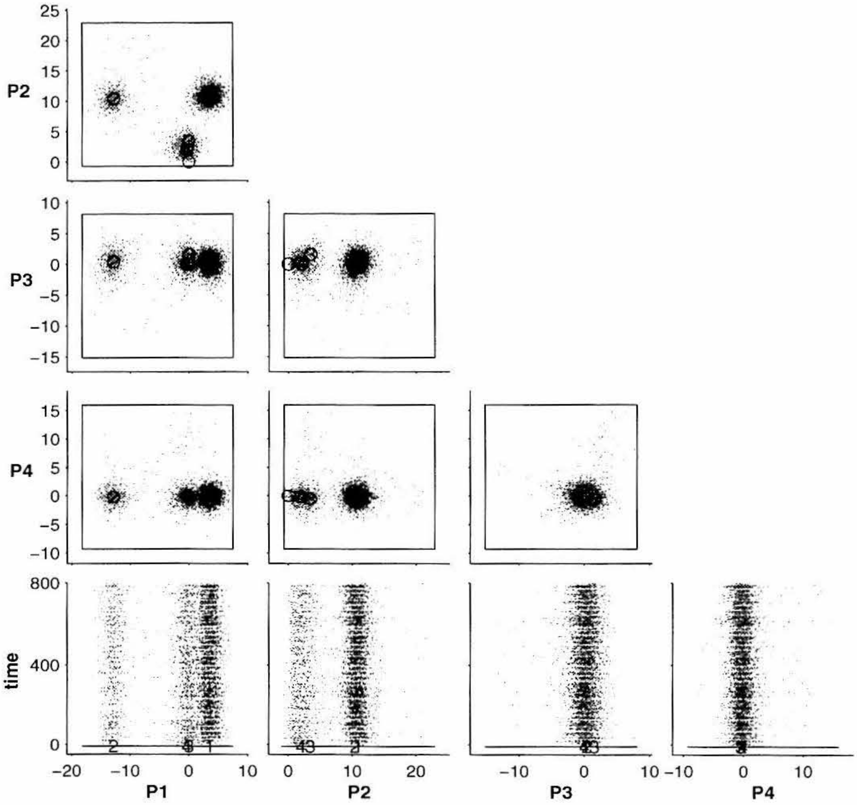

Figure 71: Sorted Events, NWrPC Space

whitened principle component space, and so when projected back to the peak space, the models appear slightly off center. This is a visualization artifact and does not reflect upon the quality of the sort.

This example has particularly well separated clusters, but even so, a number of outliers can be seen. These are due, in large part, to overlapping spikes and are normally either discarded or absorbed by the uniform background cluster.

The noise-whitened principle component projection shows a feature typical of many of our recordings, in that most of the cluster separation is available in the top two components P1 and P2. This, naturally, is by construction of these components, but we have observed that with only a small handful of exceptions, no additional information is available at P3 and above. The rare exceptions have information at P3 but never at P4. We have, by and large, sorted on the first four principle components for reasons of symmetry with the four tetrode tips; the number could be reduced to improve efficiency.

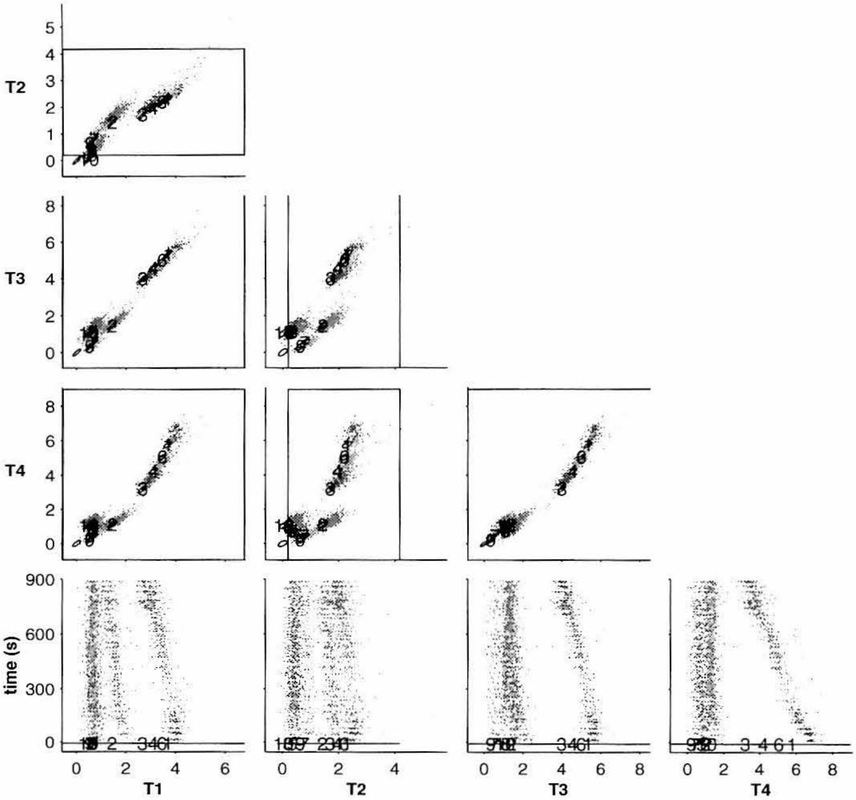

Figure 72: Clustergram Showing Drift

5 Online Processing and Drift

This sorting algorithm requires little work to be adapted to an on-line mechanism. The real limitation is one of processing power. With the computers used for development of this code, true real-time processing was out of reach, but not unreasonably so. As of this writing, processor speed has approximately doubled, which, with some judicious code optimization would probably suffice for real-time performance.

Aside from the benefits in experimental efficiency that on-line processing would bring, having a means to incrementally retrain models provides a way to compensate for drift. Drift with tetrodes under daily insertion is a problem that cannot be ignored. The fundamental source is the fine shank diameter of the electrodes, about 35 μm, which creates the problem as follows. For ease of handling, the tubes through which tetrodes are advanced must be much larger than the tetrodes (at 100 μm internal diameter, these tubes are still quite fine), allowing the electrode room to coil up ever so slightly along the length of the tube. This effective spring captures energy during electrode advancement and releases it slowly once the microdrive controlling electrode position has stopped, creating an additional advancement over time. The effect is over and above normal brain drag

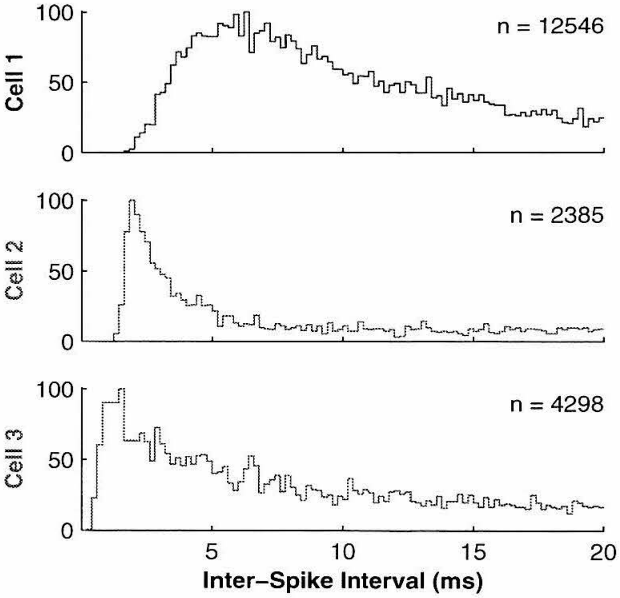

Figure 73: ISIH
The interspike interval histograms (ISIHs) for the three clusters in the example sort, normalized to 100 times their maximum values. The count in the upper right corner is the total number of intervals in each graph.

which itself often requires an hour for stabilization after electrode insertion. We have found that drift effects can be largely moderated by intentionally overshooting an area of interest by a few hundred micrometers before pulling back to the desired location.

Without employing cluster retraining, drifting clusters are tiled with multiple models which will be agglomerized in a manual post-processing step. This combining is done by setting the `cells` field of the `expt` structure; this field is a cell array, each element of which is an equivalence class of models, the exemplar for which is the first element (for cells with one model, these will be scalars). When cluster retraining has been enabled, model centers are adjusted for each trial, and tiling is not as prevalent.

Drift is detectable from tell-tale elongated clusters, as seen in Figure 72. The largest-valued cluster appears smeared-out in the upper panels of the figure; the lower panels showing peak value against time show this cluster to have a normal size over any short period of time, because the background noise is relatively constant, but changing amplitude through the recording. It is this in amplitude creates the apparent smearing. The multiple colors (or shadings) of the cluster are the result of the sorting process which has tiled the cluster with four models. The other clusters can be seen to have relatively constant amplitude; it is not unusual for lower amplitude clusters to appear fixed while higher amplitude clusters drift. It is also not unusual for drift to appear in abrupt form as if the electrode tip suddenly slipped passed some structure.

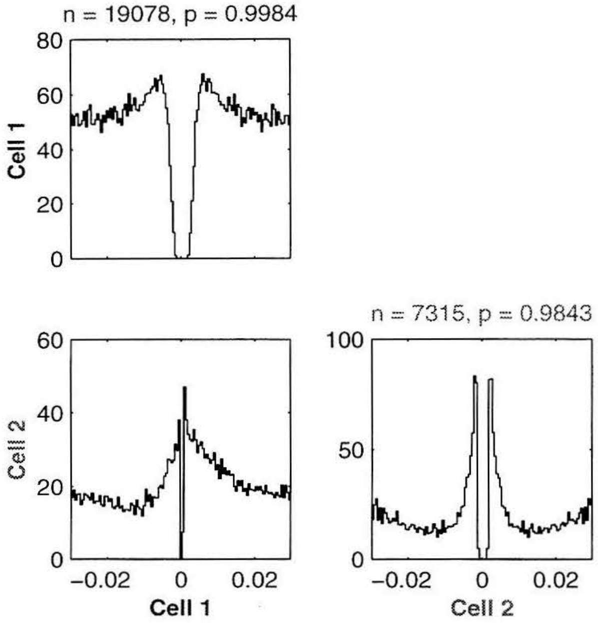

Figure 74: Auto- and Cross-Correlations

6 Cluster Verification by ISIH and Correlogram

Once a sorting has been done, the clusters are verified for source uniqueness by examining the the interspike interval histogram (ISIH) for each. The ISIH of the identified models in our example is shown in Figure 73. We look for evidence of refractory periods of at least 1 ms to verify the quality of isolation. Only two of the three detected clusters in this recording have clean refractory periods indicating spikes assigned to these cluster are from isolated neurons; the third is likely a combination of spikes from multiple cells which are all too distant to resolve into separate sources, as corroborated by its small peak height seen in the clustergrams above. The third cells' apparent, and short, refractory period is artifactual and due to a mechanism invoked during event extraction that excludes overlapped events by imposing a minimum separation between detected events.

As a preliminary evaluation of the recording for interaction between cells, the auto- and cross-correlograms for a sort are examined. The graphs for the two well-sorted clusters our example can be seen in Figure 74 spanning delays of ±25 ms; the poorly sorted cluster has been discarded. One of these cells shows evidence for bursting (Cell 2, with a pair of sharp autocorrelogram peaks at close to 1.5 ms), while the other does not. The crosscorrelogram shows an unusual interaction where Cell 2 is more likely to fire after a Cell 1 spike than before it. The notch in the center of the autocorrelograms is the refractory periods (the zero bins have been suppressed), while the notch in the center of the crosscorrelograms is due to our current inability to properly sort overlapped events.

The figures above each autocorrelogram in Figure 74 indicate the number of spikes assigned

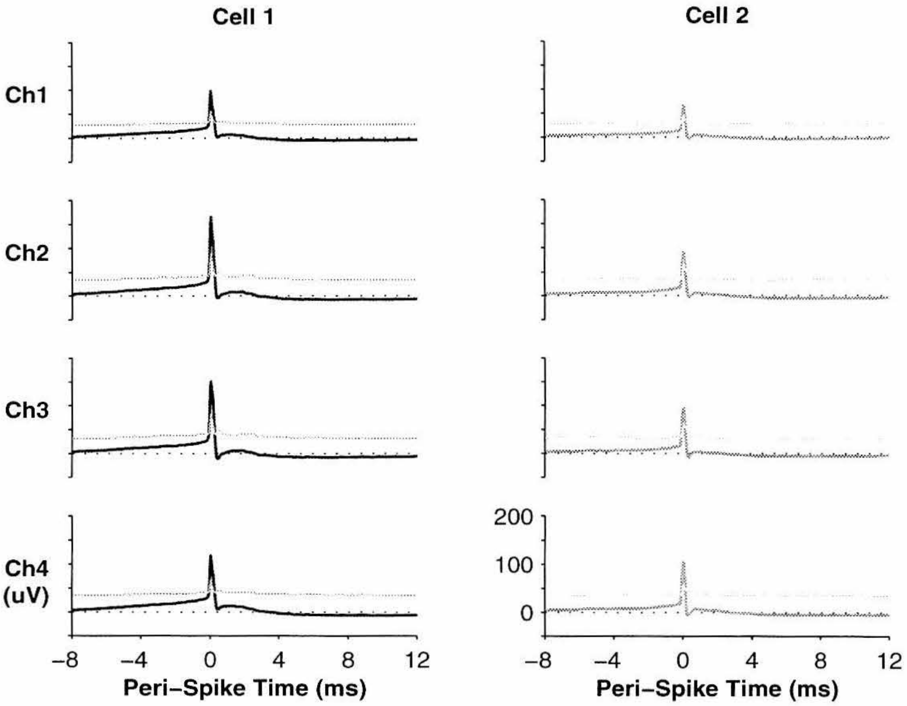

Figure 75: Spike Shapes
This figure shows the mean voltages on all four channels across 5000 events for the two cleanly isolated cells (heavy lines) and their standard deviations (fine lines). For each detected and sorted event, the original, unfiltered data was sampled for 20 ms around the spike and averaged to generate this figure. Vertical scales are all matching.

to that cluster, and the mean posterior probability among the spikes. The correlograms have been computed by propagating the posterior information from the sorting through the computation, counting each event pair (the reference spike from one cell and the observed spike from the other) not with unit weight as in traditional correlograms, but with weight of the product of the two posteriors. In addition, for this figure, only spikes with posterior probabilities of at least 0.95 were included.

7 Example Spike Shapes

The mean events for the two isolated cells are shown in Figure 75. In forming these high-n means, we examine the data from a different perspective, looking for an underlying commonality to all spikes from a given cell, and rejecting signals from all other sources. Thus the standard deviations will now show a different background noise than we have previously considered, mostly formed by the low frequency local field signals, but also from other spikes. Although the standard deviation values are high, the number of samples brings the standard error to a level approaching the width of the solid traces.

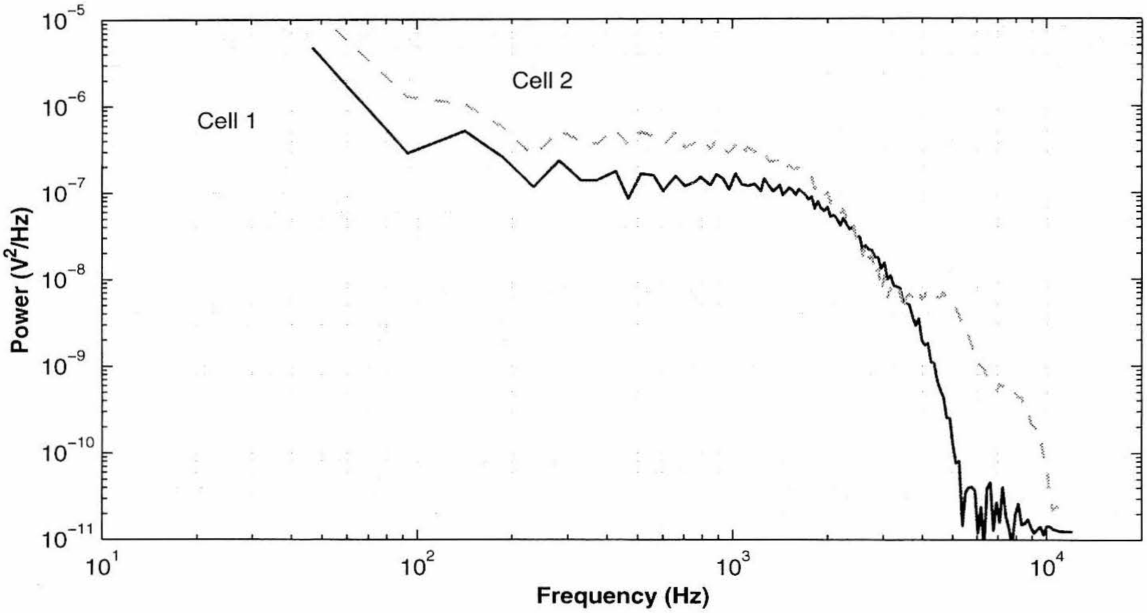

Figure 76: Spike Spectra
Spectra for the two mean waveforms from Figure 75. The four channels from each waveform were averaged and then Fourier transformed and the magnitude of the result squared to generate this figure.

These mean spike shapes show evidence of a signal surrounding each spike that spans more than 20 ms, as the voltages have not returned to zero by the edges of the figure. This runs counter to the dogma that a spike represents an isolated event lasting at most 2–3 ms. Part of this extended process will, undoubtedly be the local field, and reflect the general activity of the area, possibly including input, but part will also be from subthreshold membrane motions.

The spectra of these mean spike shapes are seen in Figure 76, and correspond well to the excess energy above the straight-line backgrounds of full-recording spectra presented in Chapter 2.

8 Summary

This chapter has reviewed the operation of our spike sorting algorithm, written in large part by Maneesh Sahani. Example sorts have been presented, along with a discussion of the evaluation tools used in verifying sort quality, including a new kind of correlogram which incorporates the posterior probabilities from the sorting process. Finally, the mean spike shapes from an example recording were presented which showed a temporally much broader effect than traditionally considered for extracellular spikes.

So far, I have not found the science
But the numbers keep on circling me
— SOUL COUGHING (So Far I Have Not Found the Science, *El Oso*, 1998)

Chapter Seven: PIST

This chapter is based on a paper originally appearing in the proceedings from CNS98 [72] called *PIST: Simultaneous Paired Intracellular and Tetrode Recordings for Evaluating the Performance of Spike Sorting Algorithms*, written in collaboration with Mike Wehr and Maneesh Sahani, describing work done in Gilles Laurent's laboratory at Caltech.

1 Abstract

Objective evaluation of spike sorting algorithms such as those used to decompose tetrode recordings into distinct spike trains requires *a priori* knowledge of the correct classification for a given recording. Intracellular recording can unambiguously assign spikes to a single neuron, and thus provide correct classification if signals from that neuron concurrently appear in a tetrode recording. Simultaneous single or paired intracellular and tetrode recordings are used here to evaluate a contemporary spike sorting algorithm for isolated as well as overlapped events. These data are also used to demonstrate that overlapping extracellular spikes combine additively, and to introduce a means for quantifying variability in action potential shape.

2 Introduction

Tetrodes [60] offer significant advantages in addressing the problem of sorting spikes from multiunit recordings into distinct trains [17, 69, 68, 23]. These bundles of fine wires have four close-packed recording sites to generate multiple views of the electrical landscape at the electrode tip. A cell which is closer to one site will produce a larger waveform there than at the other, slightly more distant recording sites. As each cell has a unique spatial position, and thus a unique pattern of distances to each group of recording sites, each neuron should, in principle, present a characteristic pattern of waveforms across the multiple channels amenable to algorithmic extraction.

Evaluating such spike sorting algorithms requires an independent verification of firing times for the collection of neurons being recorded. While synthetic data can be used for this purpose, doing so requires acceptance of assumptions about signal characteristics which may interfere with objective evaluation. An *in situ* approach, such as intracellular recording, is therefore preferable. We elected to combine dual intracellular recording with tetrode recording, allowing simultaneous verification of signals from two neurons. This provides not only the desired spike assignment confirmation, but also allows quantification of variability in extracellular action potential shape and explicit testing of the assumption of linearity in the extracellular medium during overlapped events.

3 Methods

Experiments were carried out *in vivo* on adult female locusts (*Schistocerca americana*). Animals were restrained dorsal side up, the head was immobilized with beeswax, and a watertight beeswax cup was built around the head for saline superfusion. A window was opened in the cuticle of the head capsule between the eyes, and air sacs on the anterior surface of the brain carefully removed. For stability, the esophagus was sectioned anterior to the brain, and the gut removed through a subsequently ligatured distal abdominal section. The brain was treated with protease (Sigma type, XIV), gently desheathed, and supported with a small metal platform. The head capsule was continuously superfused with oxygenated room-temperature physiological saline (in mM: 140 NaCl, 5 KCl, 5 $CaCl_2$, 4 $NaHCO_3$, 1 $MgCl_2$, 6.3 HEPES, pH 7.0).

Intracellular recordings were made using conventional sharp glass microelectrodes pulled with a horizontal puller (Sutter P-87), filled with 0.5 M KAc, for resistances of 100–300 MΩ. Intracellular recordings were done in bridge mode using an Axoclamp 2A amplifier (Axon Instruments) from the third optic lobe (lobula). Data were collected from 28 single neuron and 6 paired intracellular recordings, all with simultaneous tetrode recordings, from 7 animals.

Tetrode recordings were made using electrodes composed of 4 strands of 15 μm insulated tungsten wire (California Fine Wire) twisted at approximately 1 turn/mm and glued with common cyanoacrylate [52]. The tip was freshly cut at an acute angle using fine surgical scissors before each penetration for impedances of 0.4–0.7 MΩ at 1 kHz. Tetrodes were placed in the lobula, 50–100 μm medial to the site of intracellular penetrations (see Figure 77).

All signals were amplified, low-pass filtered at 10 kHz (8-pole analog Bessel with gain, Brown-Lee Precision), digitized at 50 kHz with 16-bit resolution (Tucker Davis Technologies), and written to compact disc. The continuous voltage recordings were digitally high-pass filtered at 250 Hz and events were obtained by threshold detection on either the tetrode or intracellular signals only. Each set of extracted events was clustered using Expectation Maximization (EM) techniques [69, 68] implemented with in-house MATLAB code.

4 Results and Discussion

Figure 78 shows a typical recording from this preparation that will be used as a running example. Each of the large spikes appears across the four tetrode channels with a characteristic amplitude signature, some with simultaneous action potentials in the intracellular channels. The tetrode recording contains signals from (at least) four identifiable cells, two of which were impaled with the intracellular electrodes. Figures 80A–B show the spike-triggered average waveforms of the six channels, as triggered on I_1 and I_2, revealing the tetrode waveform associated with action potentials from the impaled cells.

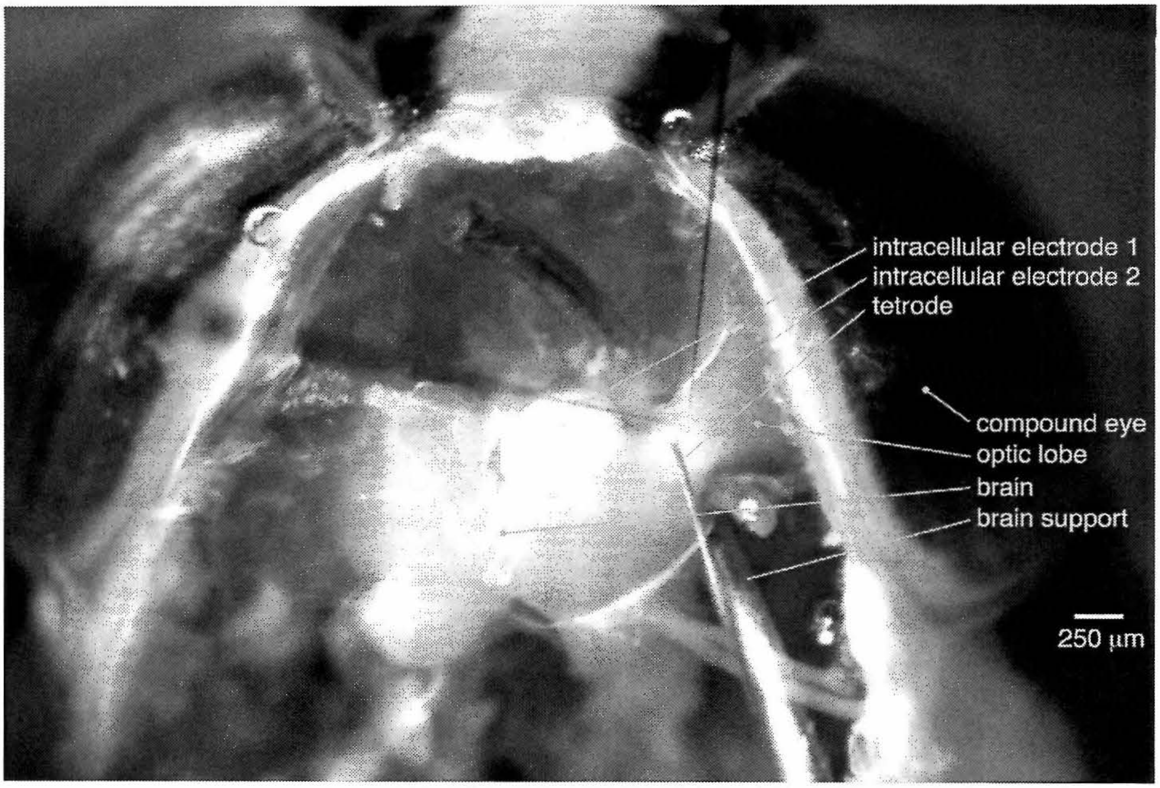

Figure 77: Photomicrograph of Locust Preparation
Photomicrograph of locust head capsule taken directly following a recording. Spontaneous activity was recorded for up to 15 minutes at a time. The retinotopic organization of the insect optic lobes, combined with the existence of large integrative neurons in the lobula, allowed simultaneous intracellular impalement of tetrode-recorded neurons at distances of up to 100 μm from the tetrode tip.

4.1 Spike Sorting Performance

The performance of a spike-sorting algorithm based on maximum likelihood techniques [69, 68] on these recordings was evaluated. The algorithm consists of two stages: (a) automated robust fitting techniques are first used to discover clusters of events, and then (b) optimal filters based on these clusters are used to decompose the signal into distinct spike trains.

During the first, clustering, stage, the purity of each cluster is of highest interest (see Figure 79). Cluster C_1 contained 862 spikes of which 831 (96.3%) came from a single cell (as verified by I_1). Only 3 additional spikes from that cell (0.4%) were misassigned to other clusters. Cluster C_2 contained 190 events with 183 (96.3%) of these from the other impaled cell (I_2). Ten spikes from this cell (5.5%) were misassigned.

At present a strong test of the second, filtering, stage has not been performed due to difficulties with suitable decorrelation (whitening) of the signal. However, a preliminary run on unwhitened data was performed to detect C_1 spikes. There were 1052 such I_1-identified events, including overlaps.

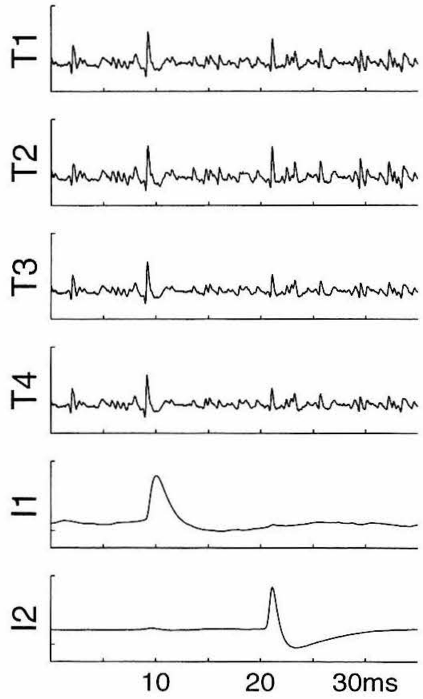

Figure 78: Example PIST Recording
Example of simultaneously recorded data. T_1–T_4 *are the four tetrode extracellular voltages;* I_1 *and* I_2 *the two intracellular voltages. Based on spike amplitudes, rising phase shapes, after-hyperpolarization amplitude, and synaptic background, the* I_1 *electrode, here, was likely impaled in a distal dendrite far from the spike initiation zone, the* I_2 *electrode in a primary neurite relatively near the spike initiation zone.*

Of these, 999 were correctly detected by C_1 filter (95.0%), however 124 additional spikes were incorrectly grouped in with these (11.8%). Such false-positives are likely to be reduced once the filtering procedure is correctly applied.

4.2 Linearity of Overlaps

To evaluate the electrical linearity of the extracellular medium, overlapping events were modeled using spike-triggered averages of non-overlapped events. Instances when the two impaled cells fired at nearly the same time were selected, and the recordings of these events were compared to the sum of M_1 and M_2 appropriately time shifted as determined from the intracellular signals (see Figure 80). For additively combined spikes, the simple sum of the models was expected to accurately predict the shape of overlap events. Actual voltage is shown as a function of predicted voltage in Figure 80E, along with a linear fit that closely matches the unity slope line, demonstrating that the additive assumption is valid across the observed voltage range.

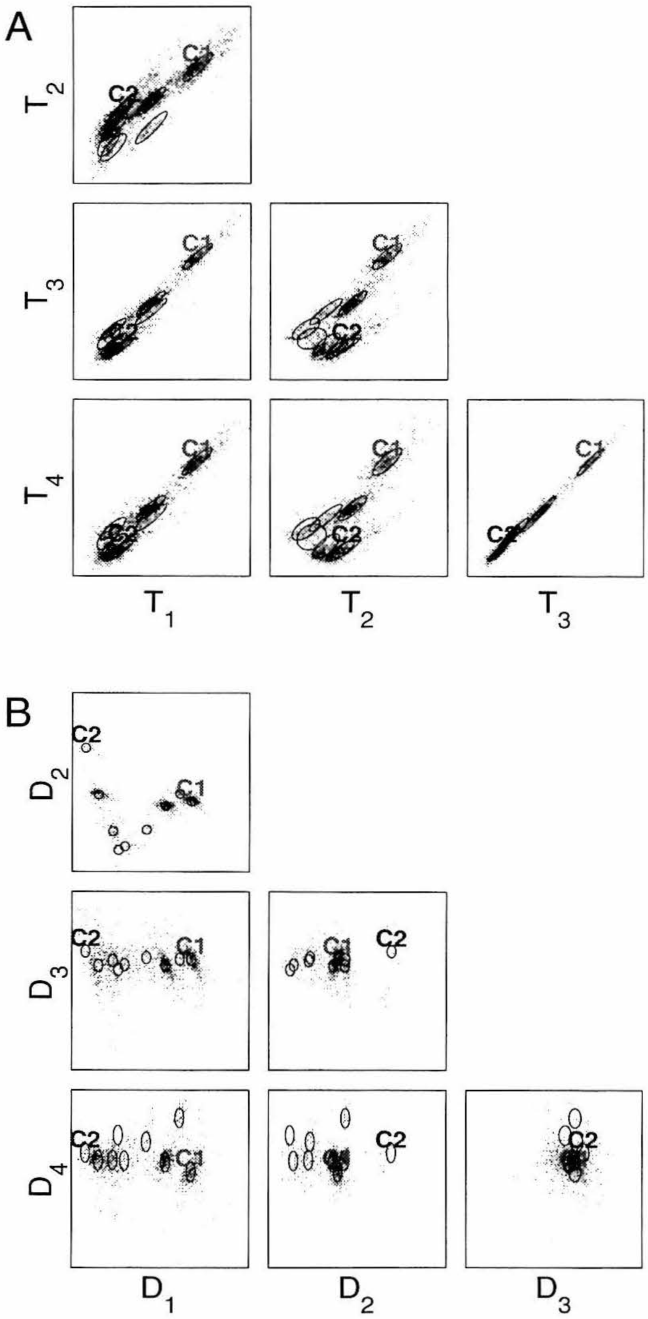

Figure 79: Clustergram of Sorting Output

Ellipses indicate fits generated from cluster analysis of the first four noise-whitened principal components. Clusters C_1 *and* C_2 *correspond to the impaled neurons from* I_1 *and* I_2*, respectively.* **A**: *Action potential peak heights for tetrode channels* T_1–T_4 *plotted against each other (axes in scaled volts).* **B**: *The same data transformed to the optimal 4D linear discriminant subspace* D_1–D_4 *(axes in arbitrary units), markedly improving the discriminability of clusters. Cluster* C_2*, difficult to isolate in* **A***, is clearly separated in* **B**.

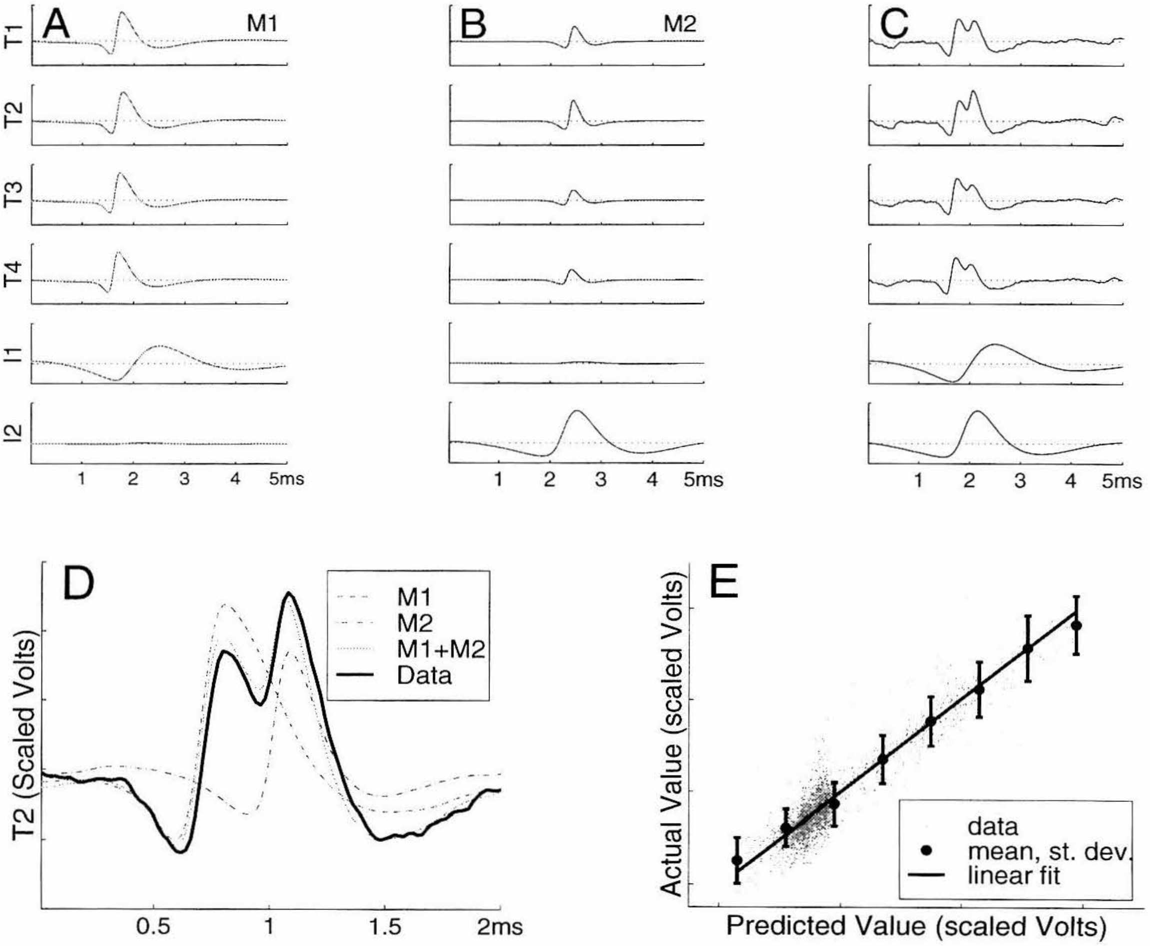

Figure 80: Spike Shapes
A, **B**: *Spike-triggered average waveforms of the six channels, triggered on each of the two intracellular channels.* M_1 *and* M_2 *are the means of* I_1 *and* I_2 *events, respectively.* I_1 *and* I_2 *have been high-pass filtered at 250 Hz for accurate event detection, resulting in waveforms uncharacteristic of standard intracellular records.* **C**: *Example overlap event.* **D**: *Expanded view of overlap* T_2, *with recorded waveform (heavy line), time-shifted* M_1 *and* M_2 *(dotted lines), and predicted combination* $M_1 + M_2$ *(fine line).* **E**: *Sample-by-sample predicted versus actual values for overlaps (dots,* $n_{event} = 88$*), with linear fit (solid line,* $y = 1.016x - 0.009$*), and standard deviations (filled symbols).*

4.3 Spike Waveform Variability

The variability of extracellular spikes was examined by computing the point-by-point covariance of all I_1-triggered tetrode events. The covariance matrix was decomposed into stationary and non-stationary parts by an EM algorithm equivalent to factor analysis in the Fourier-transformed space (see Figure 81) which, by analogy with PCA, we call Non-Stationary Components Analysis. The stationary component corresponds to additive background noise, whereas the non-stationary component corresponds to intrinsic waveform variability. The major modes of variability were computed from the non-stationary component of the covariance matrix, illustrating the ways in which spike waveforms differ from event to event.

5 Conclusions

Preliminary analysis of simultaneous paired intracellular and tetrode recordings demonstrates that one spike sorting algorithm [69, 68] performed 96% correct isolation and classification of signals from a tetrode. Further, overlapping extracellular waveforms were demonstrated to combine additively. Lastly, spike waveform variability was shown to be decomposable into stationary non-spike noise, and several modes of non-stationary variability. To obtain copies of these data sets, please contact the corresponding author (JSP).

but i know the truth
i know the whole shebang
i know the names of men they had to hang
— Soul Coughing (Unmarked Helicopters, *Songs in the Key of X*, 1996)

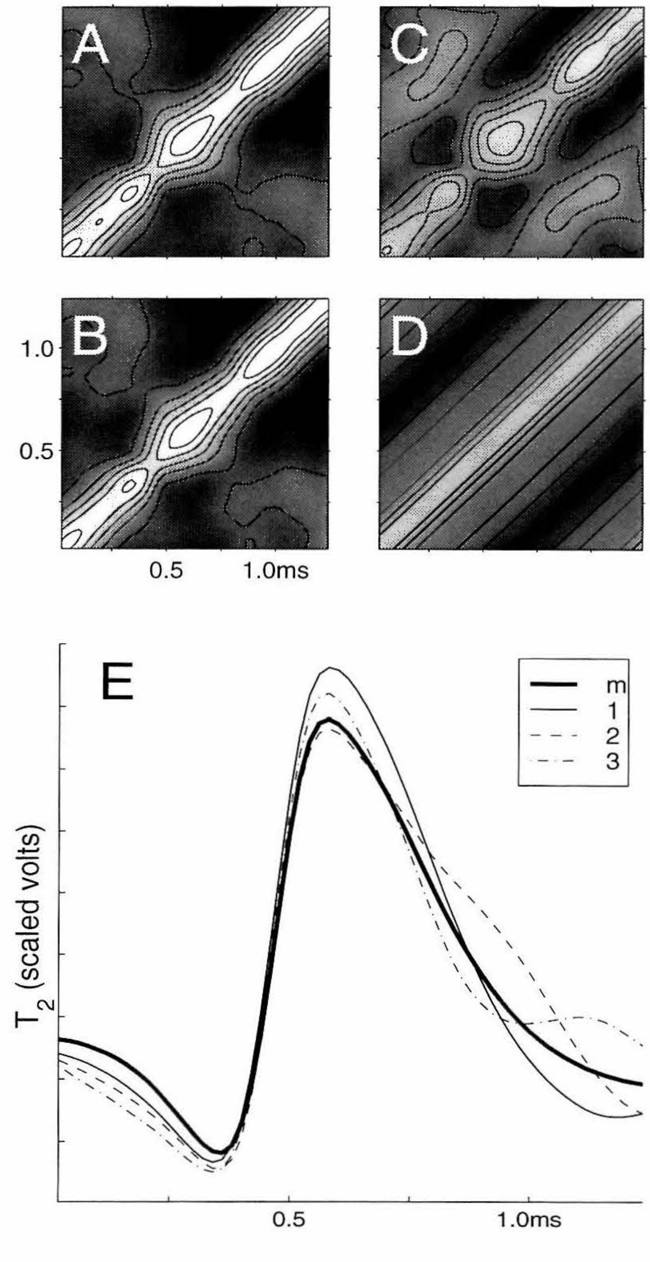

Figure 81: Covariance

Covariance decomposition of T_2 from I_1 events (**A**) *into non-stationary* (**C**) *and stationary* (**D**) *components. Reconstruction* (**B**) *from the two components closely matches the original (compare* **B** *to* **A**). **E**: *Modes of variability. The mean (heavy line) and first three modes of non-stationary variability are shown (scaled by 2). Intrinsic spike waveform variability consisted of increased amplitude (thin line), increased breadth (dashed line), or secondary peak inclusion (dot-dashed line).*

Chapter Eight: In vivo MRI Localization at 1.5T

This chapter is based on a one-page abstract titled *MRI localization of extracellular electrodes using metallic deposition at 1.5T* [51] originally appearing in the proceedings from a meeting of the International Society for Magnetic Resonance in Medicine. Since the publication of the abstract, work has been done to extend the results to additional preparations.

1 Introduction

Extracellular electrodes are widely used in neuroscience applications for electrophysiological recordings from cerebral cortex. Spatial localization of the exact recording site has traditionally relied on accurate stereotaxic positioning and histological confirmation of areas of gliosis at the end of the study. A novel technique described in rats by Fung and colleagues [19] involves electrodeposition of a small amount of iron along the electrode tract, and subsequent localization with high-field MRI. We investigated the utility of a similar approach at conventional (1.5 T) field strengths, potentially allowing neuroscience centers studying larger animal models, or without recourse to dedicated high field systems, to use more readily available conventional scanners, and thus avoid unnecessary histology.

2 Methods

The methods for this study were selected for ready availability and scientific relevance. We elected to use standard stainless steel electrodes to leave susceptible deposits, but included tetrodes as a control for potential artifacts in the primary preparation. Our studies spanned three animal preparations: bovine brain (purchased from a local butcher), laboratory rabbit, and macaque monkey. The bovine preparations were lesioned and imaged *ex vivo*, the rabbit preparations were lesioned *in vivo* and imaged both *ex vivo* and *in vivo*, and the monkey preparations were lesioned *in vivo*, and imaged *ex vivo* (after fixing) and *in vivo*.

2.1 Electrodes

Three types of electrodes were used for this study: stainless steel, tungsten, and tetrode (made of tungsten). The first two types, both traditional electrodes, were purchased from the Fred Haer Corporation (Bodoinham, Maine), the last was of the type used in collecting extracellular data found in the rest of this thesis (see chapters 2 and 3 for details).

2.2 Preparations

Our first step was to determine if the technique was viable at lower field strengths and to explore the parameter space of lesion time and current. This was performed in three *post mortem* bovine brain preparations. We then used three rabbit preparations, to verify these settings in a living preparation. A terminal monkey preparation lesioned *in vivo*, but imaged *ex vivo*, extended these results to primate preparations. And finally, a monkey preparation lesioned and imaged *in vivo* showed the technique to be a viable alternative to histological track reconstruction.

For rabbit and monkey preparations, lesions were performed while the animal was anesthetized with isoflurane at surgical levels. Under sterile conditions, 2.0 mm diameter holes were drilled through the skull through which stereotaxically mounted electrodes were positioned. Imaging was done under anesthesia (rabbit: 35 mg/kg ketamine IM, 5 mg/kg xylazine IM, 0.75 mg/kg acepromazine IM; monkey: 10 mg/kg ketamine IM, 0.5 mg/kg xylazine IM, 0.04mg/kg atropine SC).

Preparation 1: Bovine Brain

17 electrode tracts at 5mm spacing; 4 lesions per tract every 1 mm; 1–4 μA anodic current for 5–160 s.

Preparation 2: Bovine Brain

30 electrode tracts at 2mm spacing; 6–8 lesions per tract every 2 mm; 1–12 μA anodic current for 5–160 s.

Preparation 3: Bovine Brain

30 electrode tracts at 2 mm spacing; 6–8 lesions per tract every 2 mm; 1–12 μA anodic current for 5–160 s.

Preparation 4: Rabbit

8 electrode penetrations at 4 mm spacing; 5 lesions per tract every 2 mm; 4 or 8 μA anodic current for 10 or 20 s.

Preparation 5: Rabbit

11 electrode penetrations at 4 mm spacing, 8 with lesions, 3 controls without lesions; 4 lesions per tract every 2 mm; 4 or 8 μA anodic current for 10 or 20 s.

Preparation 6: Rabbit

11 electrode penetrations at 4 mm spacing, 9 with lesions, 1 control without lesions; 4 lesions per tract every 2 mm; 4 or 8 μA anodic current for 10 or 20 s.

Preparation 7: Macaque

6 electrode tracts (parietal cortex): 5 lesions per tract every 2 mm; 4 or 8 μA anodic current for

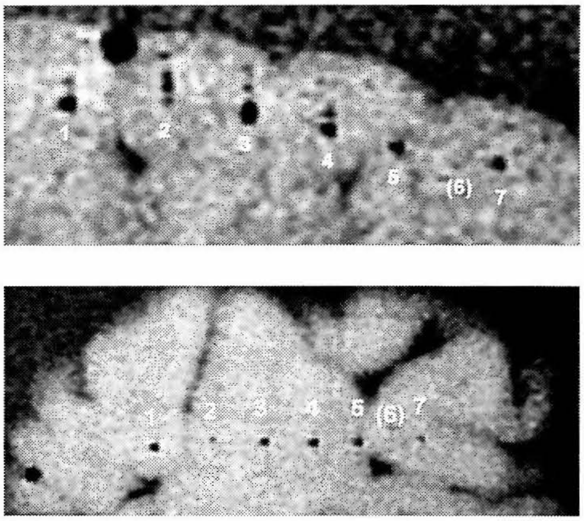

Figure 82: Preparation 1, *Ex Vivo* Bovine
Lesions in bovine brain imaged using 3D-FLASH are seen in axial (top) and sagittal sections. Only 6 of the 17 lesions were visible corresponding to electrode currents of 4μA and times of 10–160 s (currents below 4 μA and lesions of less than 10 s duration were not visualized). Table 10 summarizes the visible tracts.

Tract	1	2	3	4	5	6	7
Current (μA)	4	4	4	4	4	4	4
Time (s)	160	80	40	20	10	5	10

Table 10: Parameters for Preparation 1 Lesions
Parameter used to create the visible lesions shown in Figure 82. Additional lesions were left at similar times for lower currents, but were not successfully imaged.

20 s; MRI *ex vivo* (fixed brain). Lesions were left 1–2 hours before perfusion as an adjunct to an otherwise unrelated acute procedure.

Preparation 8: Macaque

Six electrode tracts (parietal cortex); 3 control tracts with no lesions; 3 tracts with 3 lesions each every 3 mm; 8 μA anodic current for 10 s. Tract positions abutted the intraparietal sulcus, as guided by previously done structural MRI.

2.3 Imaging

Imaging was performed at two locations, the Hoag Hospital in Newport Beach, California, and the Long Beach Hospital in Long Beach, California. Both locations had identical imaging equipment, 1.5 T Siemens VISION MR scanners (25 mT/m gradients, 300 μs rise time), and identical scanning sequences were used at the two locations.

Signals were captured using either a 35 × 17 cm flexible surface coil or a *19 cm circularly* polarized volume coil. The specimens or subjects were carefully centered in the coil, and appropriate additional loading used to insure good signal-to-noise ratios. Living subjects were sedated.

Three sequences were used. For high metallic sensitivity, a 3D-FLASH sequence ($T_R = 25$ ms, $T_E = 11$ ms, $a = 20°$, 0.4 mm and 0.5 mm isotropic resolution) was used. For intermediate sensitivity, a 3D-MPRAGE sequence ($T_R = 11.6$ ms, $T_E = 4.9$ ms, $T_I = 20$ ms, $T_D = 0$ ms, $a = 8°$, 0.8 mm and 0.5 mm isotropic resolutions) was used. As a metal-insensitive, but tissue damage-sensitive control, a 3D-RARE sequence ($T_R = 2200$ ms, $T_E = 105$ ms, $ETL = 21$, $0.6 \times 0.6 \times 0.7$ mm and 0.5 mm isotropic resolutions) was used. Total per-session scanning times were limited to 2 hours, the effective limit for single-dose sedation. In rabbits and monkeys an additional sequence 3D-FISP ($T_R = 48$ ms, $T_E = 21$ ms, $a = 15°$, 0.75 mm isotropic resolution) was also used. This is extremely sensitive to the presence of metal, but its use was limited by reduced tissue contrast and phase distortion in the final image (making exact stereotaxic location difficult).

Bovine preparations were imaged a few hours after lesioning, and the tissue was brought to room temperature before being placed in the magnet.

Rabbit preparations were imaged a few hours after lesioning. *Ex vivo* preparations were wrapped in layers of surgical draping to maintain body temperature after sacrifice. Viability of *in vivo* preparations was verified with periodic transcardiac images to confirm heart motion and pulsatile blood flow.

The first monkey preparation was initially imaged to include the whole head, 24 hours after lesioning, and then later as an extracted hemisphere. The second monkey preparation was initially imaged 24 hours after lesioning and again after 7 days. The second monkey preparation had a small cylindrical well positioned over a craniotomy as part of an unrelated experiment, but during imaging, this well was filled with gadolinium DTPA solution to verify localization.

All datasets were manipulated using in-house MATLAB code, while macaque datasets were re-oriented to stereotaxic coordinates using the AFNI software package (Medical College of Wisconsin) [13] to generate accurate coordinate references for the observed electrode tracts.

2.4 Fixing

The brain of the first monkey preparation was fixed before imaging, as part of an unrelated study. The head was perfused immediately following sacrifice with paraformaldehyde using standard methods. Tissue samples were triple bagged to prevent fixing fluid from contaminating hospital equipment during magnet work.

3 Results: Genus Bos

Bovine electrode marks were readily visualized using 3D-FLASH and 3D-MPRAGE sequences as a series

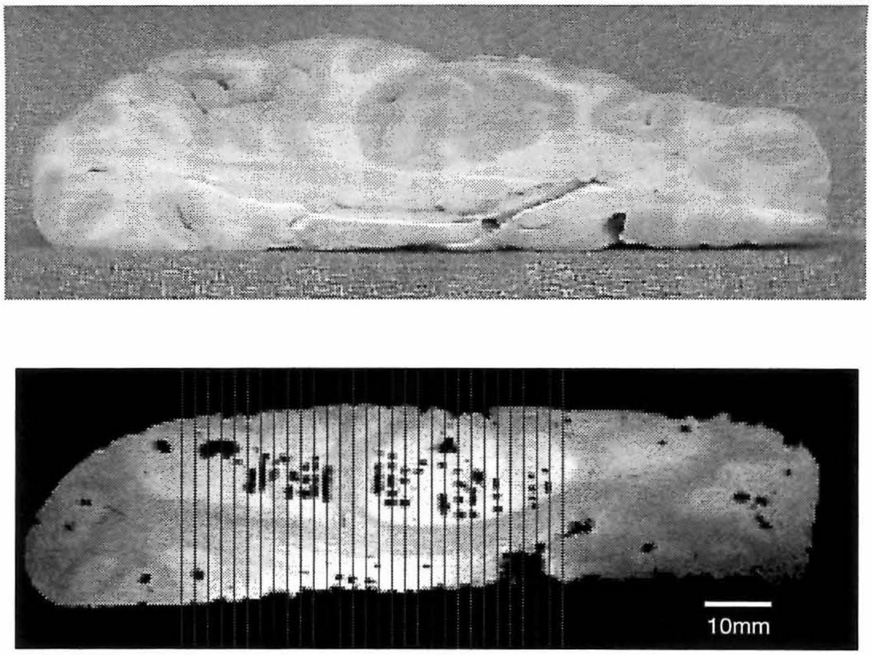

Figure 83: Preparation 2, *Ex Vivo* Bovine
Pathological specimen (top) and 3D-FLASH axial image (bottom) showing lesions are only visualized in gray matter. Vertical lines indicate axes of electrode penetrations.

of punctate areas of low signal, measuring 0.6–1.0 mm in diameter as shown in Figure 82. Lesions were not readily seen with the 3D-RARE sequence.

In the second bovine preparation, lesions from 2–12 μA were readily visualized, but only in gray matter. Figure 83 shows a series of imaged deposits (1–5 μA, 5–160 s).

4 Results: Genus Oryctolagus

All three scans, 3D-MPRAGE, 3D-FLASH, and 3D-FISP, were sensitive to both metal deposition at all currents and times, as well as to the small amount of blood that occurs in a tract as normal sequilae of electrode penetration. A large variability was observed in the detectability of lesions.

5 Results: Genus Macaca

Figure 84 shows tracts in the lateral intraparietal cortex (LIP) in fixed macaque brain. Lesions are readily visualized with 3D-FLASH (84 upper), but not with 3D-RARE (84 lower) indicating conspicuity is due to metallic deposition. Lesions were detectable in the *in vivo* preparation, as shown in Figure 85, as were penetrations without lesions, but to a much lesser extent.

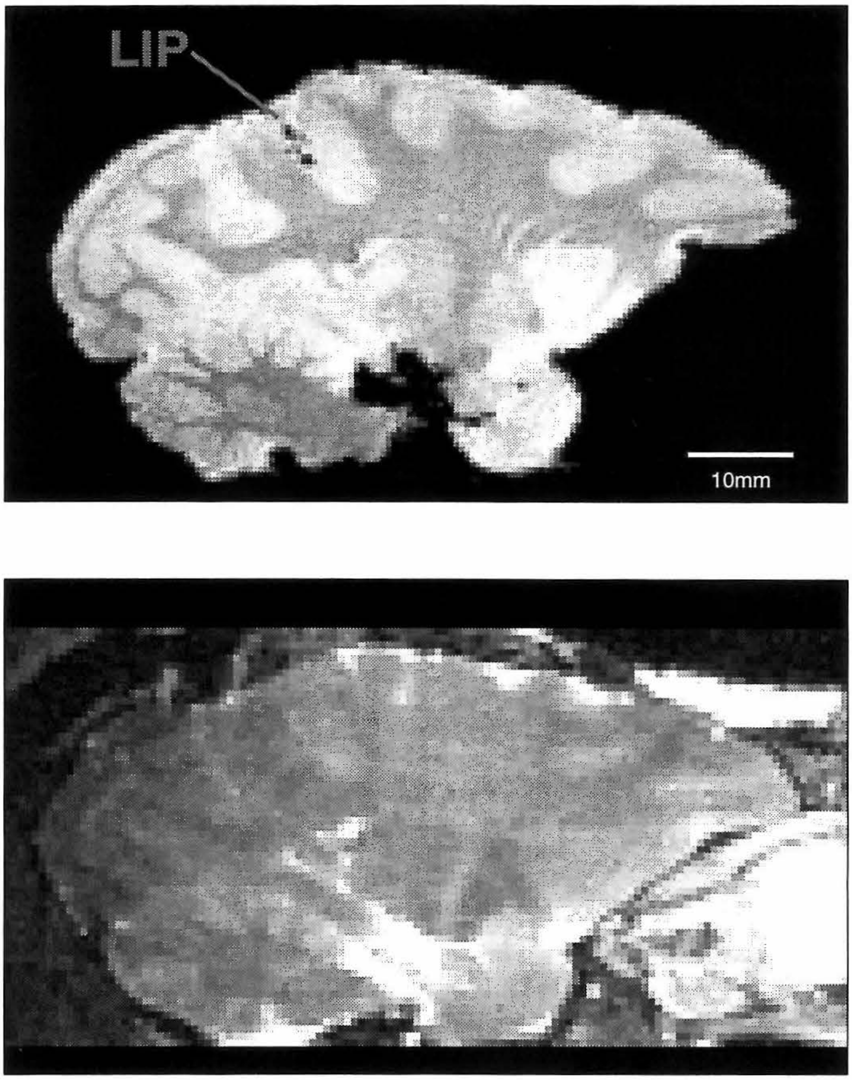

Figure 84: Preparation 7, *Ex Vivo* Monkey
Lesions in the monkey preparation are seen in under 3D-FLASH scans (top), but not 3D-RARE scans (bottom), indicating ferrous material was deposited at the lesion sites. LIP, labeled in the upper image, is a brain area on the lateral bank of the intraparietal sulcus and part of the dorsal visual stream.

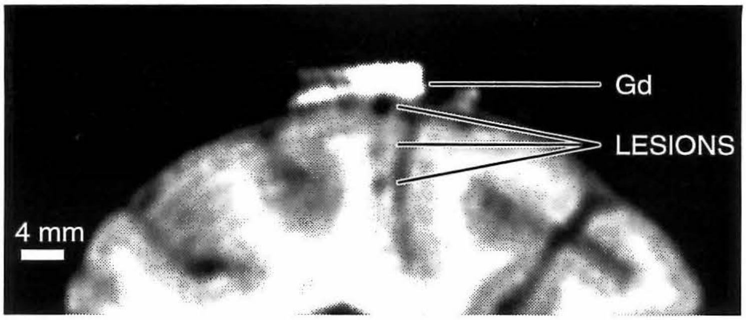

Figure 85: Preparation 8, *In Vivo* Monkey
Three iron lesions placed in the lateral bank of the intraparietal sulcus as imaged in vivo *in coronal section of a macaque brain. The bright area above the brain is a gadolinium-filled chamber used as a location marker above a previously existing craniotomy. The upper extent of the interhemispheric fissure is visible on the right, at an approximately 45 ° angle.*

6 Discussion

Electrode tracts from electrophysiological recordings can be reconstructed *in vivo* by applying unipolar currents of at least 4 μA to stainless steel electrodes for at least 20 s. Below these thresholds, electro-deposits are insufficient to reliably cause visible blooming at 1.5 T.

The failure to detect lesions with the 3D-RARE sequence confirms that the lesions do come as a result of metallic, or at least ferromagnetic, deposition at the tip of the electrode rather than local gliosis, hemorrhage, or mechanical insult.

Electrode tracts were more conspicuous with 3D-FLASH than with 3D-MPRAGE, and higher resolution was also possible. 3D-MPRAGE caused less susceptibility artifact and was more useful around areas of hemorrhage (*e.g.*, at sites of previous surgery). Electrode tract visualization was difficult if susceptibility distortion was excessive, in which case the 3D-RARE images could be coregistered with the gradient echo images and viewed simultaneously to better define the anatomy.

Additionally, lesions left in gray matter were much more conspicuous than those left in white matter. The exact mechanism for this is under investigation, and awaits histological verification of the presence or absence of iron at each lesion site.

With the first, *ex vivo*, monkey preparation as a test-bed, we developed scans which generated acceptable resolution and signal to noise within the target of 90 minutes total scanning time. This was verified in the second, *in vivo*, monkey preparation where in each of two scanning sessions a single intramuscular dose of anesthetic was sufficient to acquire 3D-MPRAGE, 3D-FISP, and 3D-RARE sequences. This time course is, we feel, realistic for *in vivo* imaging.

Based on susceptibility and signal to noise within the limits we set ourselves — readily available MRI facilities and therefore lower field strengths, reduced signal, and increased noise — we have achieved adequate resolution and lesion conspicuity to reconstruct the electrode tract.

7 Conclusions

We found that it is possible to mark electrode locations with an electrolytic lesion depositing iron ions for visualization using a conventional MR scanner within a realistic imaging time-frame. While there are reliability issues not yet resolved, unipolar current of 8 μA for 20 s generates readily visible deposits in gray matter at 1.5 T. For reasons we have not understood, deposits were not reliably visible in white matter, although in practice this is not a limitation as electrodes are normally used to record cellular activity within gray matter (not white matter). We were also able to detect localized hematomas such as occur as normal consequences of electrode penetrations.

> *I'm looking through you*
> — LENNON AND McCARTNEY (I'm Looking Through You, *Rubber Soul*, 1965)

Chapter Nine: Experimental Methods

1 Introduction

This chapter describes the experimental paradigms used to collect data presented in chapters 10and 11. There were two main behavioral tasks used, one a traditional memory saccade task, the other a fixation task used to map receptive fields. The techniques described in this chapter were used to collect data presented in this thesis from awake behaving monkeys.

2 Behavioral Monitoring

Standard techniques, already in use in our laboratory, were used to monitor animal behavior. Nearly all data were collected in the experimental area known within the laboratory as Setup 3, although a small amount of initial non-behaving data was collected in Setup 1. The equipment in Setup 3 was very nearly all built or acquired for this series of experiments, and much of it has been described in Chapter 5. The setup is now reliable and robust, and the hope is that it will last for many coming generations of students and post-doctoral scholars.

2.1 Eye Position

The animal's eye position was measured using the scleral search coil technique originally developed by Robinson [63] and improved by Judge, Richmond and Chu [30]. For each animal, a three-turn coil is surgically implanted around the sclera of one eye so that it moves with the eye, with the two output wires mechanically decoupled and routed subcutaneously to a connector on the headcap. After recovery, the animal is placed between two pairs of magnetic coils, one pair oriented vertically, one oriented transversely, driven at rationally-related frequencies (78.6 kHz and 102.4 kHz) to create two orthogonal oscillating magnetic fields of approximately uniform strength at the animal's head. Under these conditions, the signal at the headcap connector will be related to the mutual inductance between the eye coil and the driver coils which is, in turn, determined by the cross-sectional area of the eye coil as projected on the two directions. As the animal's eyes move, the horizontal and vertical angles of the eye coil determine the voltage measured at the connector, allowing a simple synchronous decoder to extract the x and y components of the animal's direction of gaze. When properly implemented, this system does not adversely affect an animal's behavior [30], and allows continuous high-speed monitoring of gaze position to reasonable accuracy ($\pm 0.5°$ near $(0°, 0°)$). As the gaze angle deviates from straight ahead, the error increases according to the error in the approximation of $sin(x)$ by x (at 30°, the error is an underestimation of approximately 5%); no compensation is made for this in our system.

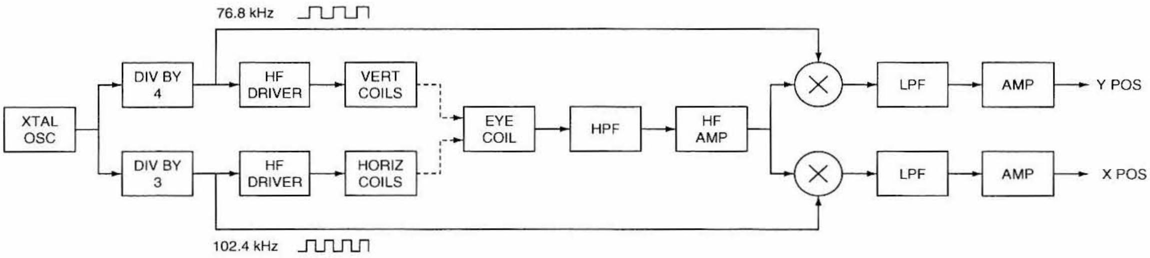

Figure 86: Eye Coil System

The driver/decoder for the eye position system was built many years ago as custom hardware, but was updated as part of this project, modifying the circuitry for additional sensitivity, better low-frequency rejection, better orthogonality of the decoded signals, higher operating frequencies, and reduction of unwanted radiation. A diagram of the system's operation is provided in Figure 86.

2.2 Posture Control

When performing the behavioral task, the animal subjects sit in specially constructed *primate chairs* adjusted to fit them comfortably, with their heads fixed by a cranially mounted post. The chairs include a chest plate which allows the seated animals to reach their mouths, but not the rest of their heads, thus preventing them from interfering with equipment at the top of their cranium. A *juice tube*, used to deliver liquid reward, is removably mounted to the chair and adjusted so that the animals can drink from the tube without straining.

An infrared camera driving a closed-circuit TV is trained on the animals during experiments to insure that they are not misbehaving, and to help assess their motivational state. The camera output is displayed on a television monitor at the top of the equipment rack, while a filtered incandescent lightbulb on axis with the camera provides infrared illumination.

2.3 Reward

The reward system was adapted [49] from the original to improve reliability and reduce operational noise. The system has a master-slave configuration with a master, electrically-actuated, valve in the equipment rack switching a compressed air source, and a slave, pneumatically-driven, valve in the experimental chamber switching a liquid reward source to a tube which runs to the animal's mouth. The electrically-actuated valve is controlled by a combination of computer input and manual override, and each pulse to the input delivers a well-metered small volume of reward liquid to the animal. Different rewards were used depending on the preferences of each animal, including Kool-Aid, Tang, apple juice, and water; the powdered drinks were mixed to approximately normal concentrations. The only drawback to this system is that as the juice tube is mechanically fixed to the primate chair, any perturbations, such as those generated by over-enthusiastic sucking, are coupled to the recording apparatus.[7]

3 Stimulus Presentation

Stimuli were presented on a screen 54 cm from the animal's eyes in the frontoparallel plane. Stimuli were points of light projected from a three-channel optical bench. Through a long series of adjustments and improvements, it was brought to a stable and reliable state. Modifications were made to reduce long-term and thermally-related drifts, decrease tangent error, eliminate light leaks, increase positional accuracy and repeatability, improve the driving interface, improve dynamic response, and, critically, eliminate the frequent failures.

The three channels are independent but identical, arranged in a linear array with parallel optical axes at a spacing of 5 inches between channels. Each channel uses a 250 W projector bulb driven by a DC power supply located outside the experimental chamber. Each channel's bulb is mounted in an aluminum frame and has a stream of cooling air aimed toward it. Light from the source is collimated by a series of lenses into a beam approximately 1 inch in diameter which passes through a high-speed mechanical shutter and continues through a pinhole aperture. Once past the aperture, the beam is deflected by a pair of orthogonally mounted galvanometer-driven front-surface mirrors and then by a large path-adjustment mirror shared by all three channels which directs the beam at the tangent screen. The light path is schematically described in Figure 87. At the projection screen, the stimuli appear as circular, elliptical, or rectangular areas of illumination approximately 1° across, depending on the aperture shape and size.

The combination of the three high-power projector bulbs in the enclosure makes high-airflow cooling a necessity. This is accomplished through a large fan with entry and exit flows though black-colored opaque ducting. The heat generated by the bulbs can adversely affect the mirror driver characteristics, so temperature-stabilized drivers are used, as described below.

3.1 Light Source

In the future, it would be desirable to update the light sources to LED lasers which will eliminate the need for the power supplies, large cooling fans, collimators, shutters, and apertures, reduce any remaining temperature-dependent effects significantly, eliminate the need for containment of stray light generated from the source, greatly reduce unintended stimulus intensity variation, allow programmed stimulus variation, and likely eliminate the need for periodic bulb replacement. The one drawback to this plan for using a pre-focused light source, such as an LED laser, is that arbitrary apertures are not readily available.

---

[7] Regrettably this was not discovered until after our data were collected, and although it affects only the local field potentials during the inter-trial interval, it should likely be addressed.

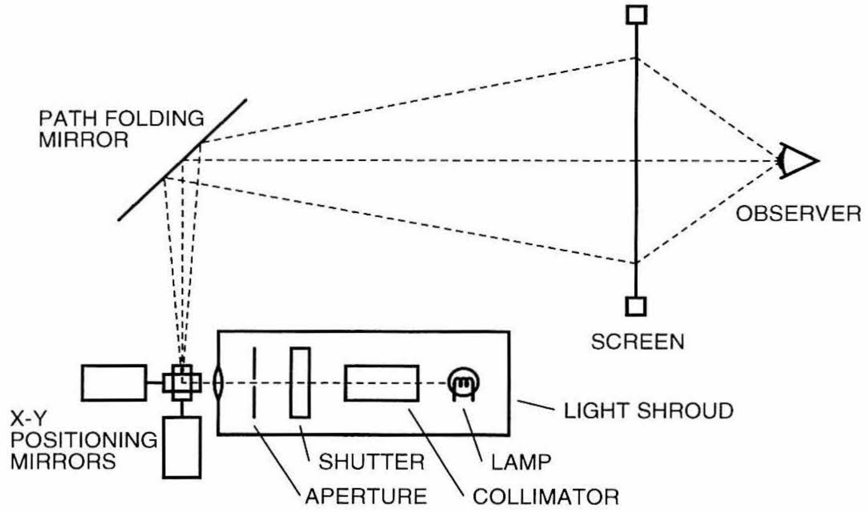

Figure 87: Light Path
The light path, schematically depicted in side view, with one of three channels shown. Light travels from the high-intensity bulb through a collimator, computer-controlled shutter, pinhole aperture, to orthogonally-mounted x-y positioning mirrors, a path-folding mirror, and the rear of the translucent projection screen. Rotation of the vertically mounted mirror about it's long axis controls vertical displacement of the stimulus. The dashed lines show examples of three different vertical displacements of the beam. A shroud encloses the optical bench to eliminate any extraneous leaks.

3.2 Light Path

For each channel, light travels from an incandescent bulb along an optical rail to a shutter, aperture and the two controlled mirrors, a large path-folding mirror, and finally to the tangent screen, as shown in Figure 87 and analyzed in Figure 88. The optical bench, from the light source to just before the x-y mirrors is surrounded by an aluminum shroud painted black and sealed against light leaks that provides paths for cooling air flow. The light source, for each channel, is a high-powered 250 W, 24 V halogen incandescent bulb (EHJ24V-250W) with a remote DC power supply. The bulb is mounted in a four-sided housing, missing rear and top, with a cooling air stream flowing from the rear. The front of the housing holds the first of a series of collimating lenses, most of which are in a cylindrical holder mounted on the optical rail. A high-speed shutter (Uniblitz model 225L2A0T5) mounts on the rail and the far end of the collimator slips into the shutter. A few centimeters behind the shutter, a frame is placed on the rail which accepts aperture cards. After this there is an additional lens, through which the beam exits the optical bench shroud. The beam then strikes the two x-y mirrors, exits vertically, and is deflected back towards the tangent screen by the large path-folding mirror.

3.3 Mirror Control

Each controllable mirror, two per channel, is driven by a precision DC motor with positional feedback via a PID (proportional-integral-derivative) controller. A command voltage delivered to the controller drives the mirror to a calibrated position. The accuracy is high, but the linear error components are corrected as described below. The assemblies are heated by the controllers to 45°C with thermal blankets to mitigate any temperature-dependent effects.

3.4 From Stimuli to Positions

Stimuli are specified in a spherical coordinate system based on optical angle as viewed from a reference position where the animal's eyes are placed. The geometry illustrated below was used to derive an expression linking desired optical positions to equivalent system input voltages. In operation, behavioral control programs will express target location in degrees of visual angle; these figures are transformed, checked for validity, and delivered to a digital-to-analog converter to create a command voltage which is sent to the mirror motor controllers. This mapping from degrees of visual angle to A/D voltages is described below.

3.5 Corrections and Calibration

Because of construction and manufacturing tolerances, there is a certain amount of discrepancy between commanded and actual dot position. This is corrected to linear errors with a simple procedure: in uncompensated condition, the system is commanded to $(0°, 0°)$ and $(20°, 20°)$ and the actual positions read. These values are transcribed to a centrally located computer file, and the correction terms are then automatically computed during normal operation. When calibrated in this way, the error is empirically limited to less than 1°.

The calibration correction is performed before the command signal is converted from stimulus coordinates to command coordinates and sent to the mirror controllers. The transformation and correction are computed with the formulas

$$x' = (x - z_x)g_x \tag{15}$$

$$y' = (y - z_y)g_y \tag{16}$$

$$\phi_x = \arctan\left((d_{m,s}\tan x' + k_p d_{p,p})/d_{p,s}\right) \tag{17}$$

$$\phi_y = \arctan\left(d_{m,s}\tan y'/d_{p,s}\right), \tag{18}$$

where (x, y) is the intended position, (x', y') is the position after linear gain and offset correction, (ϕ_x, ϕ_y) is the transformed position, g_x and g_y are the gain corrections and will be close to 1, z_x and z_y are the offset corrections and will be close to 0, $d_{m,s} = 33 \pm 0.2$ inches is the distance from the monkey to the screen, $d_{p,s} = 80 \pm 0.4$ inches is the distance from the projector to the screen, $d_{p,p} = 5 \pm 0.01$ inches is the distance from projector to projector, and k_p is -1, 0, or 1 for projector number 1, 2 or 3, respectively. The values for g_x, g_y, z_x, and z_y are computed during calibration. The transform is graphically depicted in Figure 88.

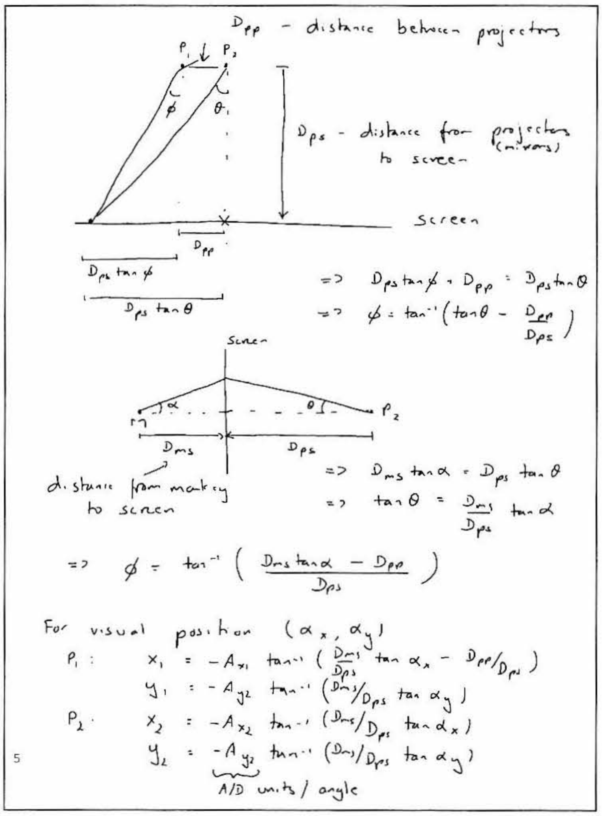

D_{pp} - distance between projectors

D_{ps} - distance from projectors (mirrors) to screen

$$\Rightarrow \quad D_{ps}\tan\phi + D_{pp} = D_{ps}\tan\theta$$

$$\Rightarrow \quad \phi = \tan^{-1}\left(\tan\theta - \frac{D_{pp}}{D_{ps}}\right)$$

distance from monkey to screen

$$\Rightarrow \quad D_{ms}\tan\alpha = D_{ps}\tan\theta$$

$$\Rightarrow \quad \tan\theta = \frac{D_{ms}}{D_{ps}}\tan\alpha$$

$$\Rightarrow \quad \phi = \tan^{-1}\left(\frac{D_{ms}\tan\alpha - D_{pp}}{D_{ps}}\right)$$

For visual position (α_x, α_y)

P_1: $x_1 = -A_{x_1} \tan^{-1}\left(\frac{D_{ms}}{D_{ps}} \tan\alpha_x - D_{pp}/D_{ps}\right)$

$y_1 = -A_{y_2} \tan^{-1}\left(D_{ms}/D_{ps} \tan\alpha_y\right)$

P_2: $x_2 = -A_{x_2} \tan^{-1}\left(D_{ms}/D_{ps} \tan\alpha_x\right)$

5 $y_2 = \underbrace{-A_{y_2}}_{\text{A/D units / angle}} \tan^{-1}\left(D_{ms}/D_{ps} \tan\alpha_y\right)$

Figure 88: Sahani Lab Notebook 1, Page 5
Upper diagram is a plan view of beam paths with optical bench at top and screen at bottom. Lower diagram is a side view, left-right reversed from as depicted in Figure 87.

A more sophisticated correction would, for example, involve measurements at an array of points (say every 5 or 10°) so that a two-variable second degree polynomial could be fitted to the uncorrected positions. It was empirically determined that this was not necessary for reasonable accuracy, and further, that the calibration was stable for time spans of months. However, should additional accuracy and stability be required, the first steps to be taken should be to improve the corrections to compensate for mechanical irregularities, and to change from the current incandescent light sources to solid-state sources, eliminating the multi-hundred Watt heat source the projector bulbs currently present.

4 Neural Access

Recordings were made with tetrodes through recording chambers which were mounted above craniotomies in acrylic islands fixed to the skull. Tetrodes were constructed using 13 μm diameter insulated tungsten wire, twisted, and lightly fused to hold their shape, as extensively described in Chapter 3. Daily insertions were made using a guide-tube-like carrier tube mounted in a standard hydraulic microdrive and positioned over the area of interest. Accepted and approved sterile surgical techniques were used to form the craniotomy and implant the recording chamber and acrylic island with headpost, also known as a *headcap*. Recording chambers were cleaned before each use, and regularly when not in data-collection phases of the experiment, and antibiotics applied to control infection. Periodic removal of hyperplastic tissue was performed as necessary.

Carrier tubes were advanced through the dura by manual control on the microdrive, after which tetrodes were advanced under remote hydraulic control until signals from neurons could be detected on an oscilloscope and on audio monitors. When a well-isolated set of task-responsive neurons was found, recordings were made with equipment described below. Downward in-tissue motion of tetrodes for all non-recording penetrations was made in 1 μm steps at 3 steps per second. Upward motion was often increased to 5 μm steps at the same step rate. Adjustments made during fine isolation varied in speed, but were always made in 1 μm steps.

5 Signal Conditioning and Capture

Neural signals were amplified and conditioned using a mixture of custom and off-the-shelf hardware. A custom low-noise headstage amplifier and secondary amplifier were designed and constructed, as presented in Chapter 5. The headstage has a gain of 100, high-pass rolloff at 0.15 Hz, and low-pass rolloff at 15 kHz. The secondary amplifier has adjustable gain, additional high-pass rolloff at 0.15 Hz, and a low-pass rolloff at 25 kHz. The system gain was nominally set to 20,000, but for some recordings was lowered to 10,000 or even 5,000 to optimize recording representation and prevent amplifier saturation.

Amplified signals passed through anti-alias filtering at 6.4 kHz (earlier recordings) or 10 kHz (later recordings) and digitized to 16 bits at 12.8 kHz (earlier recordings) or 24 kHz (later recordings) at ±10 V full scale. Before filtering, at least one channel of each recording, usually channel 1, was routed through a switch which, under the command of the behavioral program, alternated between the neural signal and ground to mark the start of each trial and allow synchronization between behavioral events and the neural recording.

5.1 Spike Train Analysis

Continuously recorded voltages were converted to spike trains using custom software written in cooperation with, and in large part by, Maneesh Sahani [69, 67]. This algorithm is covered in more detail in Chapter 6, and will be reviewed only briefly here. For a full treatment, the reader is encouraged to seek out the cited references.

To start the conversion process, the signal was digitally high-pass filtered at 300 Hz, following which a two-phase analysis was used: models were generated for spike shapes, and then these models were applied as filters to create spike trains. A two-phase process was used in part because the length of the recordings precluded simultaneous analysis of every spike, and in part because an architecture readily extensible to on-line analysis was being developed.

In the first phase, optimal models were generated. A sampling of 100 s of each recording was made typically using 1 s segments evenly spaced through the entire recording (recordings were hundreds of seconds long). This data stream was whitened, both in time and across channels, and events crossing 5 times the covariance for at least two samples were detected. These events corresponded to neural spikes. At each threshold crossing, a 2 ms segment of the unwhitened stream was extracted, upsampled by a factor of 4 using Fourier reconstruction, shifted so that the center of mass above 2 times the covariance was aligned to a reference position in the segment, typically 3/8 of the way through the segment, and decimated back to the original sampling rate. Extracted events were then nominally whitened, and projected to a principal component space, of which the top four dimensions were typically retained. The samples were rotated to form the optimal discriminant for that space, and a probabilistic clustering algorithm used to segregate the events into model classes.

In the second phase, these models were used to generate optimal filters which were applied to the continuous data stream to generate spike trains. Recordings which exhibited drift, or contained spikes with a high degree of variability would often have multiple filters for each cell, these parsed spike trains were merged. Each spike train was verified to be from a unique cell by insuring that refractory period was visible on the inter-spike interval histogram. Spike trains were then combined with behavioral data to form trial-based data structures to assist later analysis.

6 Experimental Tasks

Two tasks were used for this work, a memory saccade task, and a visual receptive-field mapping task. For both tasks, the animals were trained until nominal performance was in the 70–90% range before recording commenced, and ongoing performance levels used as an indicator of motivational state. One monkey performed these tasks willingly without water deprivation, presumably for the flavor of the reward, while others required water deprivation for effective motivation.

Behavioral monitoring and stimulus generation was under the control of the computer program Hydra, written by Maneesh Sahani, fragments of which were presented above. This real-time experimental control system allows sophisticated tasks to be easily programmed in a high-level language.

Because the neural signal is noisy, to extract useful information about a site, we need multiple repeated sampling of a given experimental condition. Each sampling is called a *trial*, and a collection of sequentially performed trials is called a *block*. Normally, a block contains a small set of different trial conditions which are pseudorandomly interleaved, and arranged such that approximately equal numbers of each condition are present. During analysis, trials belonging to given conditions are collected and the data pooled in an effort to increase the signal to noise ratio.

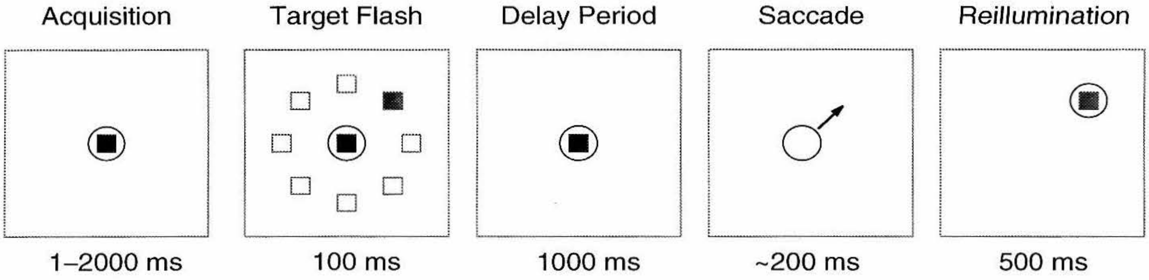

Figure 89: Memory Saccade
The experimental paradigm.

6.1 Memory Saccade Paradigm

The *memory saccade task* is a direct adaptation of the experimental paradigm originally developed by Hikosaka and Wurtz [27] and extensively used by many other investigators. Our version is graphically depicted in Figure 89. The animal sits in a primate chair, with head fixed. The chair is placed in the driver eye coils so that the animal faces a tangent screen, onto which points of light are back-projected. The entire apparatus is housed in a shielded room painted black and maintained dark for the duration of the experiment. Care was taken to reduce ambient light levels so that the only visual stimuli for the animal are the ones presented experimentally. In addition, as the apparatus in the room requires substantial cooling, external noise sources are masked by the sound from the cooling fans.

Experiments were conducted in blocks of 100–150 trials. Typically each block of trials corresponded to a single recording location, but, for particularly interesting sites, several blocks would be collected. For a small set of recordings, more than 150 trials were collected.

Each trial proceeds in time as follows. A fixation point is illuminated near the animal's central gaze position. While this point is present, the animal is required to foveate it. A pseudo-random interval after fixation, a target point is briefly flashed in one of eight locations evenly spaced around a circle, typically 10° in radius, centered around the fixation point. The animal was required to remember the location of this target through the following delay, or *memory period*, the end of which was indicated by the fixation point being extinguished. At this cue, the animal was required to saccade to the remembered location of the target. If successful, the target was re-illuminated for 500 ms, often triggering a corrective saccade, and the animal received a drop of liquid reward.

For most experiments, the targets were arranged every 45° with positions starting directly rightwards of the fixation point and passing counter-clockwise. For other experiments, the phase around the target circle was adjusted to stimulate the center of a measured receptive field. For others, the number of targets was reduced to four or even two to maximize probing of receptive field maxima. And for yet others, the number of radii was increased to 3, 4 or 5, evenly spanning a typical range of 4–16°. These variations will be discussed as appropriate in subsequent chapters.

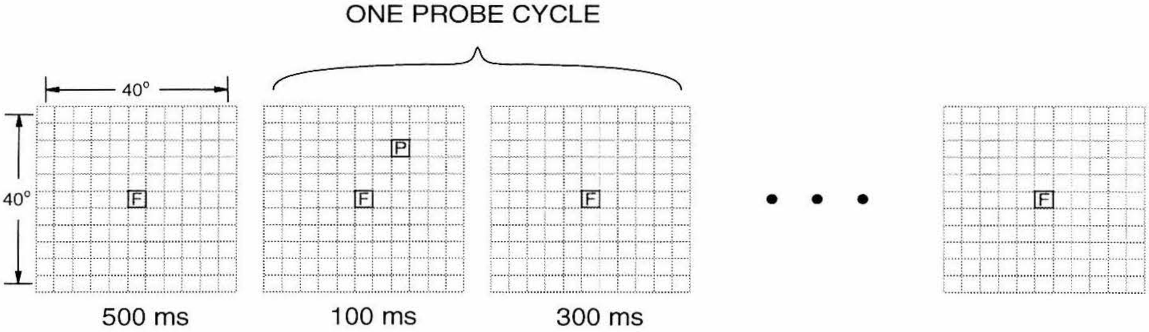

Figure 90: Receptive Field Mapping
The experimental paradigm. Each trial consisted of a series of 10 probes while the animal fixated a central point, where each probe was a 100 ms flash of light at a pseudorandomly selected location followed by 300 ms of recovery. See the main text for additional detail. This figure, and the similar maps presented below, is slightly misleading because it implies that the stimuli were 4° across, when, in fact, they were only 1° across, but placed on a 4° grid.

6.2 Receptive Field Mapping

The *receptive field task* was designed to quickly and simultaneously assess the location of multiple cells' visual receptive fields. An 11-by-11 matrix spanning $\pm 20^\circ$ along both x and y axes is mapped by repeated probes with 100 ms flashes. At the start of each trial, the animal is required to foveate a central fixation light which is illuminated for the duration of the trial. 500 ms after fixation, there is a series of 10 probes to pseudo-random locations; each probe consists of a 100 ms flash of a 1° target followed by 300 ms of recovery. A typical trial takes approximately 6 or 7 seconds, including the inter-trial interval. Fixation is required to be maintained within a 4° diameter circle around the fixation point for the duration of each trial. A block of 66 successful trials was collected while neural data were recorded for off-line analysis that was performed immediately following acquisition as the results were used to guide subsequent experiments. The task is schematically depicted in Figure 90.

It was empirically discovered that reducing the inter-probe recovery time to 100 ms does not substantially change the deduced receptive fields, although it does reduce the signal-to-noise ratio, and therefore the certainty of centers and widths. As there are apparently interesting dynamics in the responses to such stimuli out to at least 300 ms post-flash, we used the longer inter-probe interval.

7 Summary

This chapter reviewed the experimental methods used to collect neural and behavioral data. Stimuli were generated with a three-channel optical bench, eye position signals captured using a standard scleral eye coil technique, and neural signals filtered and recorded using custom hardware. Neural signals were converted to spike trains using a state-of-the-art clustering method. Two experimental paradigms were used to collect the data. The first, and more extensively used, was the memory saccade task which requires that the animal remember the position of a flashed target for approximately one second. The second, and less extensively used, was the visual receptive field mapping task which was used to more accurately measure the spatial extent of neural responses to flashed targets.

> *Monkey see, monkey do, I don't know why.*
> — Nirvana (Stay Away, *Nevermind*, 1991)

Chapter Ten: Firing Rate Based Analyses

This chapter examines recordings made in the basic paradigm explained in previous chapters using tools based on first-order information, that is, firing rate. The next chapter will use tools based on second-order information such as correlation and coherence. Some of the work in this chapter has already appeared in a conference proceedings [55], but is presented in extended form here.

1 Introduction

Multiple single unit extracellular recordings were made using tetrodes in macaque posterior parietal cortex while the animal was performing a visual memory saccade task. Recordings were made over a 2×2 mm area at both superficial and deep locations in two hemispheres from two different animals. Spikes were separated using an Expectation-Maximization algorithm to fit clusters to the distributions of either spike peak height or the first four principle components of spike shape. 47 sites were selected for analysis based on task response and clarity of separation, yielding 143 total neurons with a mode of 3 cells per site. The response criteria and stereotaxic location used were consistent with identifying neurons within the lateral intraparietal area (LIP).

For cells within the set of selected sites, responses to the task were further categorized based upon spatial characteristics (preferred target and direction) and temporal characteristics (time of maximal response). Neighboring cells were found to be likely to have similar tuned direction (63% within one octant), but not as likely to have similar temporal characteristics. We take this as evidence that area LIP is heterogeneous at the local level.

Temporal responses (peri-stimulus time histograms) were analyzed with principle component analysis, which developed modes of variation similar to our informal notions of the different phases of responses in LIP, and two clustering algorithms, which produced very poor results. The lack of clusterability is taken to imply a continuum of response characteristics, rather than a discrete set.

This chapter also presents results from mapping techniques to quickly assess the visual receptive fields of neurons, with a comparison made to the fields recorded with a standard memory saccade task in LIP. Additionally, results from a memory saccade task with denser spatial sampling are presented which provide support for the hypothesis that a potential confound for our negative result on clusterability was from using fixed stimuli locations that did not necessarily present targets at the center of neuronal receptive fields.

Finally, the spectrum of each cell was computed and examined for discriminability from that of a variable-rate Poisson simulation. Contrary to our previous reports, many cells exhibited evidence of oscillation[8].

2 Experimental Methods

As the methods have been extensively introduced and detailed in preceding chapters, they are summarized here for convenience.

We trained two rhesus monkeys (*Macaca mulatta*) to perform the memory saccade task, monitoring eye position using the scleral coil technique [30]. The task, performed in a darkened room, requires the animal to fixate a centrally-presented light while a peripheral target is flashed in one of eight equally-spaced positions on a 10° radius circle about the fixation point. The animal must remember the location of the target as long as the fixation light is illuminated, delaying the instructed saccade until the fixation light is extinguished. If the saccade is performed accurately, the target is briefly reilluminated, and the animal rewarded with a drop of juice. LIP neural response to this task is characterized by three phases: a sensory response to the target flash, an elevated baseline during the delay period, and a perisaccadic burst during the cued eye motion, often combined with a second sensory response during target reillumination. Not all cells in LIP exhibit all phases of response.

To collect neural responses, tetrodes [60] were repeatedly inserted into cortical tissue near the intraparietal sulcus through a chronically maintained craniotomy over the lateral intraparietal area (stereotaxic coordinates 6 mm anterior, 12 mm lateral; locations in one animal (L) have been histologically verified as being in LIP, those from the second (C) have not yet been verified), using techniques previously described [56]. The four tetrode voltages were filtered, digitized, and streamed to digital media while the animal performed the task. These recordings were then analyzed off-line using statistical techniques described by Sahani, *et al.* [67, 69, 68], and the resulting spike trains examined for spatial tuning and temporal response profile as described in the following sections. Neurons were considered to be *neighbors* if they were recorded simultaneously at the same site.

Penetrations were made between 3000 and 10000 μm down from the putative cortical surface, to arbitrary depths (often every 500 μm along a single penetration). At the sites where an on-line isolation could be made using any single channel of the four tetrode signals with traditional (single-channel) equipment, the animal was run on a block of 80–120 trials, and the tetrode voltages recorded for later analysis. Sites reported below were further selected for clarity of signal, multiplicity of cells, and either brisk response to the task during the memory phase, or proximity to a previously recorded site with a strong response.

---

[8] The term *oscillation* is often mistakenly used in the Neuroscience literature where a much softer term is intended. Our usage carries the connotation from Physics of a process with mechanical regularity.

3 Results

Results will be presented in five subsections. The first will analyze the behavioral performance of the two animals on the task to insure that the animal can perform the task in both cases; the second will give a general description of LIP neural responses to the memory saccade task; the third will discuss responses across the spatial dimension, looking for target specificity; the fourth will discuss responses across the temporal dimension, looking for task-phase specificity; finally, the fifth will examine the power spectra of recorded cells for oscillation.

3.1 Behavior

To insure that recordings can be analyzed in a cohesive manner, we must verify that the two animals were performing reasonably and comparably. To this end, various durations of behaviorally-relevant events were computed, such as the actual duration of the stimuli, the saccadic delay, and so forth. The results are shown in histograph form in Figure 91, with the means and standard deviations in Table 11. The values for target duration, memory period, and target reillumination delay are primarily machine-determined, and thus computing them from the collected data should merely serve to convince ourselves that the behavioral programs were working as designed. It does point out, however, one experimental oversight: the memory period was intended to be a single, fixed value, however it is evident that at some point a change was made in the controlling code such that we have two values for memory period, 950 ms, and 1200 ms. This means that in computing results based on aligned data, we must not go beyond 950 ms into the memory period when aligning on the target flash, if we wish to combine results from the two conditions (which broadly correspond to the two animals). Similarly, since the target reillumination latency takes on two values (170 ms and 190 ms) for one animal and only one (170 ms) for the other, we must be careful when examining data aligned near those events.

The two means for saccade latency, 235 and 200 ms, support the hypothesis that despite any anticipation they might have had, our animals were waiting for the fixation point offset before initiating a saccade. The smoothness of the overall distributions neither show evidence for a separate express saccade distribution, nor for temporal quantization of latency. However, these data have not been examined for individual recordings, thus evidence for quantization may be hidden by the pooling process. If the data are binned to sufficiently fine resolution, a strong quantization effect does appear with a 4 ms period, but this is believed to be an experimental artifact due to technical details of the means used to collect the eye position data and represents an inherent granularity of temporal measurement rather than behavioral limitation.

Although the two saccade latency means are significantly different (two-tailed t-test, $p < 0.001$), they differ by less than 15%, and the two distributions substantially overlap, thus we feel confident that the behaviorally-linked observations for each animal can be validly compared. We

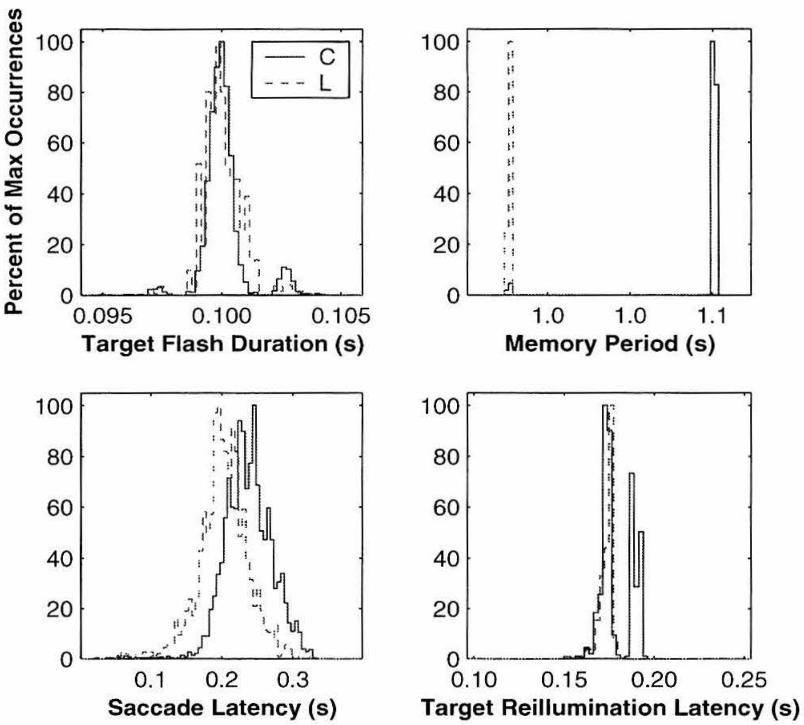

Figure 91: Behavior
Distribution of time periods from trial. The target flash duration is expected to have very small distribution, as it is a machine-generated period. The memory period is similar, however, as is evident from the distribution, two different values were used. The saccadic latency is a measure of the response time of the animals, and shows little evidence of express saccades. The target reillumination delay is a combination of saccadic flight-time and a programmed delay after target acquisition.

	Monkey C	*Monkey L*
Target Flash	0.100 ± 0.001	0.100 ± 0.001
Memory Period	1.190 ± 0.047	0.950 ± 0.001
Saccade Latency	0.235 ± 0.034	0.200 ± 0.035
Target Reillum	0.179 ± 0.009	0.171 ± 0.004
N trials	3300	1400

Table 11: Mean Durations
The mean values (in seconds) for the distributions in Figure 91. Note that the standard deviation for Monkey C's memory period is unusually high because the distribution is strongly bimodal. Numbers of trials have been rounded to the nearest 100.

will therefore examine each set of recordings from the two animals at the same time through the remainder of this chapter.

3.2 General Response to Task

As electrophysiologists working in area LIP, we are accustomed to intuitively decomposing neural reactions to the memory saccade into the three phases of visual, memory, and perisaccadic responses. An example response for this task is given in Figure 92 which shows each of these phases; this recording will serve as a running example through much of the chapter. This full plot contains a spatial map of the responses to each target, with subgraphs placed according to the equivalent target location. Each subgraph contains two panels: the left spans trial times from slightly after fixation onset to slightly before fixation offset, including the target flash, to which times from each trial are aligned; the right spans trial times from slightly before the saccade to well after, aligned to the saccade. The upper part of each subgraph holds rastergrams[9], peri-stimulus time histographs (PSTH)[10], eye position traces[11], and behavioral markers. The behavioral markers for this experiment are the illumination of the fixation point (dark, or blue, band from −0.5 to 1 on the left subpanels), the target flash (lighter, or red, band near 0 on the left subpanels, drawn on top of the underlying fixation point indicator), the saccade (vertical line at 0 on the right subpanels), and the target re-illumination (lighter, or red, band to the right of 0 on the right subpanels).

Examination of the target responses shows that of the four cells in this recording, three respond preferentially when targets are presented in the upper left position, and remain active through the memory period despite there being no further stimulation once the target flash has finished. This is traditionally taken as evidence for maintenance of a short-term memory for the target [7, 40]. Also, closer examination of the response to the preferred target (as will be done later in the chapter) shows that each of the three target-selective cells shows slightly different shaped PSTHs, including an increase directly after the target flash, known as the *visual burst*; sustained *memory activity* between the target flash and fixation offset which can be increasing, decreasing, or relatively constant; and a *perisaccadic burst* surrounding the saccade (this effect is weak in our example; a later example shows a stronger effect). Thus, we have tuning in space (target selectivity) and in time (structured PSTH). These two aspects will be considered in turn.

3.3 Spatial Response Profiles

Many of the recorded cells exhibited spatial tuning over target location. The tuning for our running example is given in figures 93 and 94, in spatial map, and polar plot forms, respectively, where three of the four cells can be seen to show target selectivity. To quantify the response, the average

---

[9] A *rastergram* has time running on the horizontal axis, trial number on the vertical axis, and dots placed for each detected spike. Often, as here, the axes are unlabeled.

[10] A *PSTH* shows mean firing rate versus time. The firing rate is taken as a mean over all trials, and the diagram can be though of as depicting the running density of the rastergram representation.

[11] *Eye position traces* show measured direction of gaze in x and y components versus time. For the x component, positive displacement is rightwards motion; for the y component, positive displacement is upwards motion. The data for successive trials are overplotted.

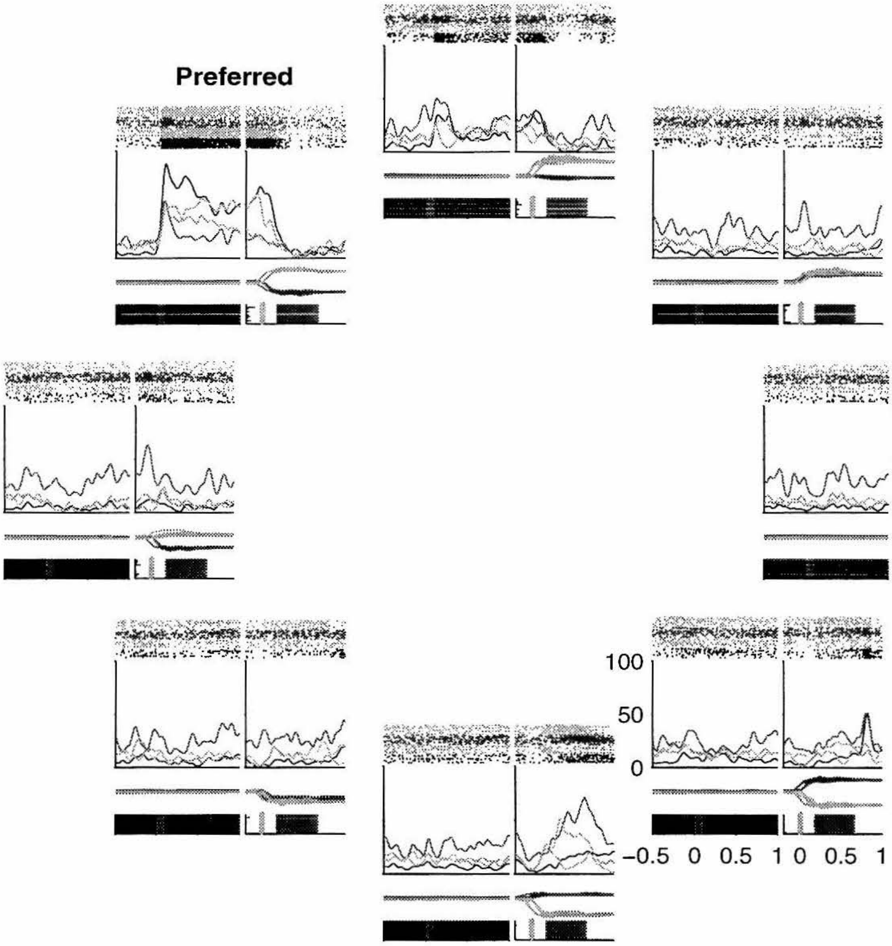

Figure 92: Example LIP Response to Memory Saccade (Full)
Four simultaneously recorded cells from area LIP (lmem2810) shown in a spatial map. Each subplot corresponds to the response to a target presented in the equivalent position on the stimulus screen. Within each subplot are rastergrams, PSTHs, *eye position traces, and behavioral markers (additional detail can be found later in the chapter), with a common horizontal time axis spanning time in trial. Three of the four cells recorded at this position responded to the task, all preferring the target position to the upper left, each with a slightly different response. Two of the cells also have a visual response to the target re-illumination towards the end of the trial for targets at the bottom.*

number of spikes for an experimentally relevant period of all successful trials is counted for each target direction. Consistent with previous reports [21, 8], we find many of the cells recorded express target selectivity within the task.

To determine tuning significance, we used a bootstrap method based on the responses from target flash through 200 ms after the saccade. This computation showed 102 of 143 cells (71%) were significantly tuned ($p < 0.05$); in the example curves shown in Figure 91, three of the four responses

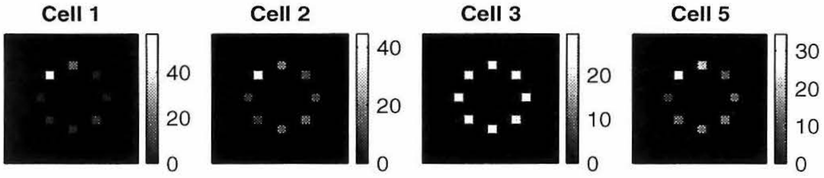

Figure 93: Example LIP Response to Memory Saccade (Map)
The same four simultaneously recorded cells from area LIP (lmem2810). Each square is placed at the position corresponding to the associated target, and colored according to the cell's memory-period activity for that target. Cell numbers are determined by the automated spike sorting algorithm; in this case spikes from Cell 4 were not cleanly isolated, and therefore were discarded.

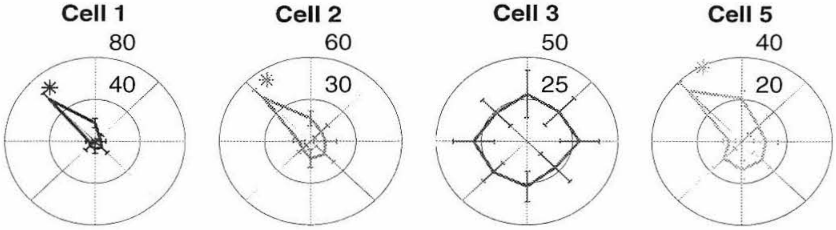

Figure 94: Example LIP Response to Memory Saccade (Polar)
Data from our running example, presented in polar form. For each cell, the task response is computed and plotted so that the radial excursion of the curve represents the magnitude of the response for the corresponding target direction, and error bars denote the standard deviation of the response. Preferred directions, computed as described in the main text for cells with significant tuning, are marked with asterisks.

are significant. The analysis assumed an excitatory response, and classified the infrequent inhibitory responses as untuned, thus figures which take these additional cells into account will be higher.

As the spatial sampling of our stimuli is reasonably coarse due to experimentally-driven constraints, we interpolated between the eight sampled positions to estimate the *preferred direction* of each cell to higher resolution by computing the vector average of the eight response vectors constructed from the origin to each of the data points on the polar graph. While there are many methods for this, including fitting circular Gaussians, cosine functions, and the like, we opted for a direct, quickly computed function. The computed preferred directions are indicated by asterisks in Figure 94, and will be used in the next section.

Pairwise Spatial Tuning

As introduced above, we extracted the preferred direction for each cell, which then allowed comparisons to be made between these values for neighboring cells (those which were simultaneously detected in a single recording). The scattergram of preferred direction among pairs of neighboring cells is shown in Figure 95. The data cluster around the unity-slope line, suggesting a tendency

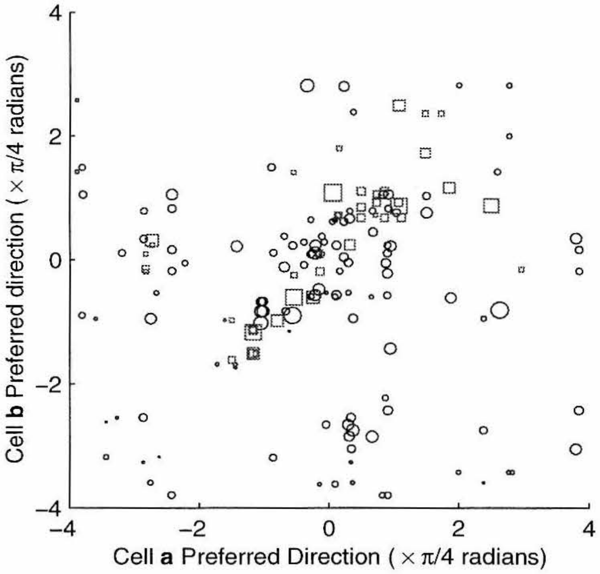

Figure 95: Scattergram of Preferred Directions
Scattergram of preferred directions for neighboring pairs of cells (a, b). *The preferred direction was computed for each cell by taking the vector average of the response for each target direction, and pairwise plotted for neighboring cells. The area of each symbol (circles, Monkey C; squares, Monkey L) is proportional to the product of the magnitude of the maximum responses in each neighboring pair.*

for neighboring cells to have similar preferred direction. Considering previous single-unit findings on the patchwork nature of LIP [21], outliers on this figure may represent recordings made when the tetrode was situated at the transition between two groups of neurons with distinct clustered preferred directions.

To quantify the similarity of tuning direction, we compute the histogram of differences between preferred directions for neighboring cells, as shown in Figure 96. The majority of pairs of neighboring neurons (63% between the two animals) have preferred directions within $\pi/4$ radians, corresponding to the cluster of points around the unity-slope line in Figure 95. Points on the graph are also concentrated in the $-\pi/2$ to $\pi/2$ range because of the tendency for LIP receptive fields to lie in the contralateral hemifield.

Spatial Mapping: Visual Receptive Fields

Having seen clear evidence for spatial selectivity in a one-dimensional probe, we ask to what extent this selectivity spreads. Because of the difficulty in collecting sufficiently many trials to probe a large number of positions, two avenues were pursued. The first was to map only the visual response, which can be done relatively quickly, taking about 1/3 the time (5–6 minutes) of a normal 8-target recording, and examining the similarity between visual and memory response fields. The second was to improve recording stability to the point where three or four times the number of trials to be taken

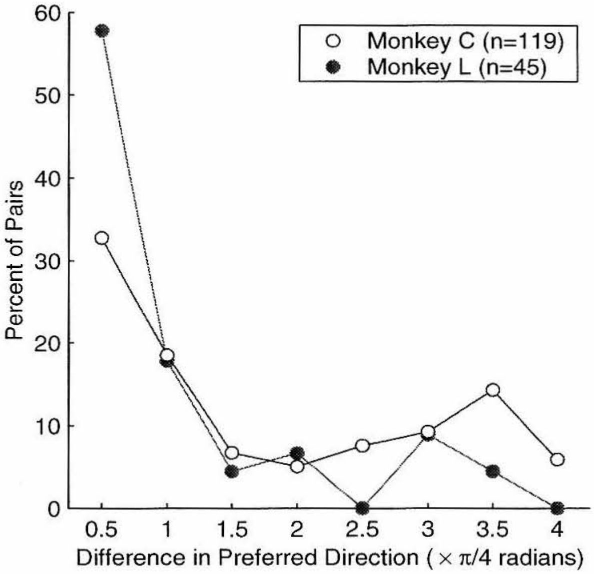

Figure 96: Differences in Preferred Directions
The differences in neighboring preferred target direction has been binned to half octants (abscissa values are bin upper thresholds). The difference in direction, a circular quantity, has a maximum value of π radians; we have then binned these values to eight ranges. The majority of pairs of cells have similar preferred direction, as indicated by the large leftmost bins.

at a given site without losing isolation, allowing multiple radii to be used for each of the normal 8 target positions. While both avenues are promising, neither has had sufficiently many recordings to allow meaningful quantitative analysis. Therefore, we present one example of each as anecdotal evidence that the techniques are promising.

The first example uses the previously presented mapping paradigm to produce visual response maps for a 40° × 40° array about a central fixation point. The map for an example recording with two cells is given on the left in Figure 97, and the 8-target memory-saccade response on the right. There is good agreement between the two: we see that the two receptive fields are very similar in this static presentation despite the visual field spanning only 300 ms of response, while the memory saccade field spans 1.3 s.

Examining equivalent portions of each task—300 ms starting after the target presentation—at a finer time scale shows very similar responses, as seen in Figure 98, except for a higher firing rate in the memory saccade case. This may be due to the higher saliency of a single target, differing levels of motivation as the animal is trained to ignore mapping task targets but remember memory saccade targets, or differing levels of anticipation since the memory saccade target does not appear for at least a full second after fixation is acquired. Both maps show short-lived waves of activity for both cells; the primary volley in cell 1, 50–100 ms after target flash initiation shows an on-response latency consistent with published values [40], however, the secondary volley, at 250–300 ms post

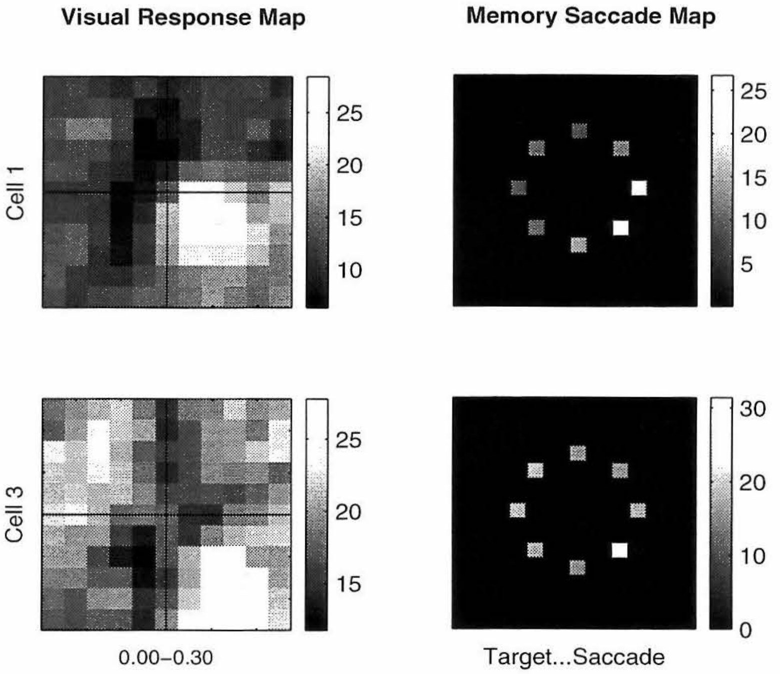

Figure 97: Visual Map and Memory Saccade Map
Results from two recordings made at the same location, the visual receptive field on the left for two cells, the memory saccade response on the right for the same two cells. The visual fields form a good predictor for response on the memory saccade task.

flash initiation will be 150 ms after the target flash has finished, and thus has latency which is inconsistent with it being a visual response to the target offset.

Spatial Mapping: Multiple Radii

Although the visual receptive field mapping shows excellent results, the interactions between different portions of the receptive fields of multiple cells can only be observed with a more spatially fine-grained measurement in the full memory saccade task. To this end, we made a small number of multiple-radii recordings, the results from one of which are presented in figures 99, 100, and 101. While substantially more difficult to perform as the stability criteria become much more stringent, the recordings provide results which illustrate the extent and shape of the receptive fields. For clarity of presentation, we show responses from only one of the three cells from the recording. The other two cells exhibited visual, and post-saccadic responses, respectively.

The three figures represent the data in three different ways. The first, in Figure 99, is the summed response for the experimentally relevant epoch of target flash through to saccade completion, shown in a spatial map akin to the previous maps. The second, in Figure 100, uses the same representation for a sliding-window analysis with 200 ms non-overlapped increments. The dynamics in the response are now clearly evident, as is the different relative sizes of the visual and memory

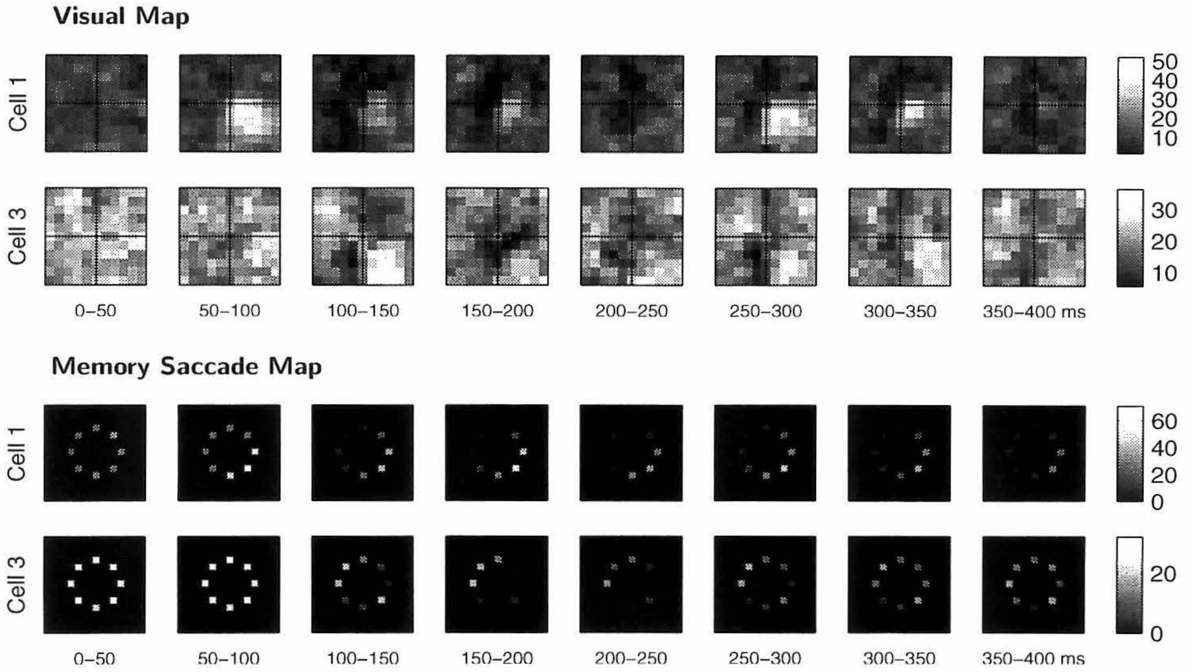

Figure 98: Fine Time Scale Responses to the Two Tasks
The target was on for times 0–100 ms in this figure, and each panel represents a 50 ms incremental slice of the mean response. The upper figures have been smoothed to reduce the effects of noise.

receptive fields. The third, in Figure 101, shows the data in spatially-arrayed PSTH form. The cell has a preferred target location of 4° eccentricity, directly rightwards of the fixation point, for which it vigorously fires through the entire memory period. The response is not limited to this target, although it starts relatively more broadly spread at the target flash, and narrows through the memory period.

3.4 Temporal Response Profiles

A more detailed example of memory saccade task response at a single site where three neurons were identifiable is shown in Figure 102 (this figure is taken from our first example, although one cell has been dropped). The graph depicts the PSTH for the target location which evoked the strongest response for this site. The two largest valued profiles (solid and dot-dashed lines) show typical LIP responses for the memory saccade task: a brief burst of activity in response to the target flash followed by a period of sustained memory activity, and a lesser perisaccadic burst. The third cell (dashed line) has a slower initial response, perhaps even lacking a distinct visual burst, followed by a decaying memory response and no perisaccadic burst. In the remainder of this section, we will attempt to quantify these various shapes of responses.

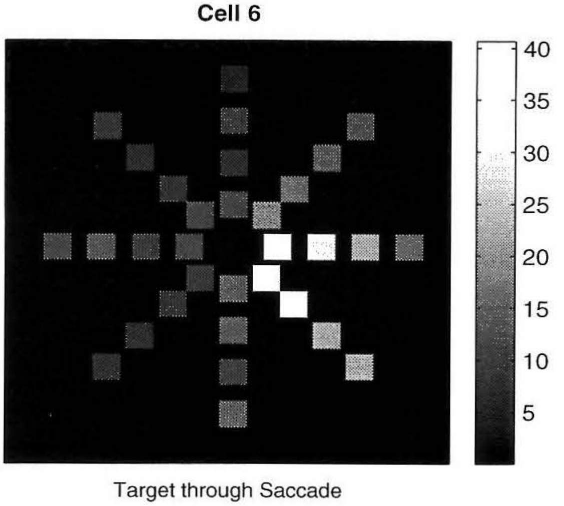

Figure 99: Multi-Radius Map
Mapping of number of spikes from target flash through saccade completion for four different radii (4°, 8°, 12°, and 16°) for each of the eight target positions. Compare with figures 100 and 101.

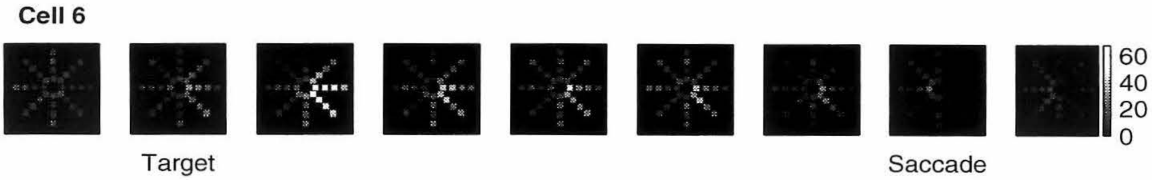

Figure 100: Multi-Radius Map
Mapping broken out over time slices. Each panel is 200 ms of time, starting 200 ms before the target flash, through 200 ms after the saccade. The size of the response field can be seen to shrink over time, remaining centered at the point of initial maximal response.

Pairwise Temporal Tuning

The response profile of each cell was computed by measuring firing rate to the preferred target during eight arbitrary but experimentally relevant epochs spanning the period from the target flash through completion of the saccade, each lasting 200 ms. Plotting the response versus direction for each epoch, we find that neighboring cells vary considerably in the time of maximal response. More striking than our last example, the simultaneously recorded pair shown in Figure 103, the upper cell fires most vigorously in the perisaccadic epoch while the lower cell fires most vigorously in the visual epoch, despite their preferred directions (during the epochs of highest response) matching quite well.

We compared the distance in time between periods of maximal response for neighboring cells and created the histogram, as shown in Figure 104. The temporal response for each cell was computed in the eight epochs as described above, and the time of maximal response used as an index of temporal tuning. The temporal distance between tunings from two cells in neighboring

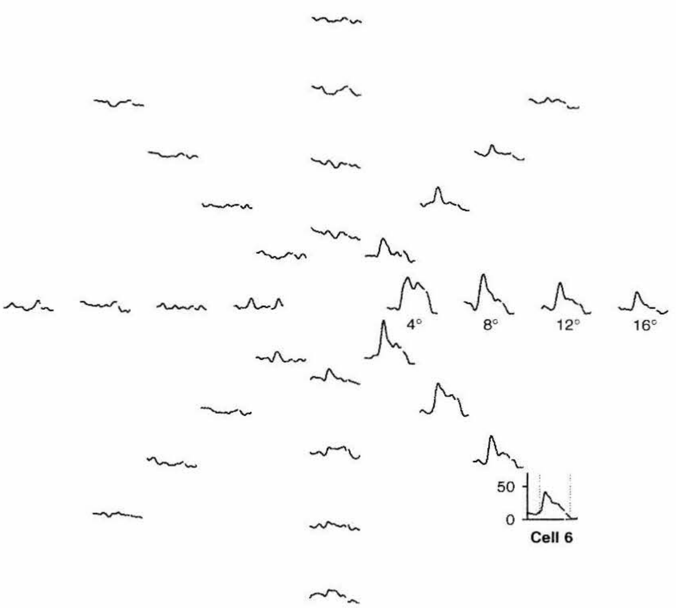

Figure 101: Multi-Radius Map
Alternate presentation of data from Figure 100, of PSTH in spatial map. By comparing across stimulus positions, the extent and dynamic shape of the response field can be seen centered at the 4° rightward stimulus position, larger at first, and shrinking over time. The dotted vertical lines indicate the target flash and saccade, and are alignment points for the two subgraphs in each plot.

pairs was then computed as the difference in tuning index. While not a rigorous metric due to the often non-unimodal characteristics of the responses, the wide distribution of the result suggests that neighboring cells often have different computational tasks. The metric used here intended merely to assist a gross determination of similarity among response types: as not even 20% of the pairs peak in the same epoch, and only about 25% of pairs peak in adjacent and once-removed epochs, we conclude that the response characteristics have a broad tendency to be different among neighboring cells. This issue will be re-examined below.

PSTH Principal Components

To further characterize the temporal variability in responses, we computed the PSTH for the preferred target for each cell, binned at 200 ms, starting 400 ms before the target flash, and continuing for 200 ms after the saccade. The mean of the first two elements was subtracted from each of these 9-tuple responses to remove the background activity, and the principle components of the results computed. The first four principle components are shown in Figure 105, and correspond to memory without visual, extended visual without post-saccadic, short visual with post-saccadic, and perisaccadic activity.

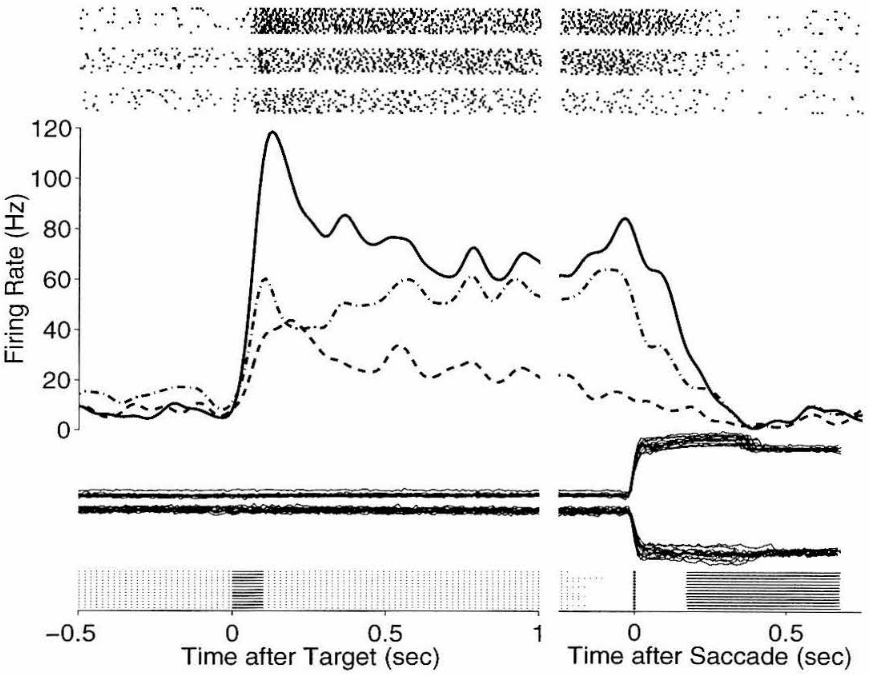

Figure 102: Example Neural Response

The left half of the figure is for a portion of the experimental task aligned to target presentation; the right half is aligned to the saccade. There is a small but varying amount of overlap between the two halves. At the top of the diagram are three rastergrams, separated out by cell; the corresponding rows from each group of rasters are from the same trial. In the middle are the peristimulus time histograms (PSTHs) for the same three cells, binned to 10 ms and smoothed to three bins. The solid line corresponds to the uppermost rastergram, the dash-dotted line to the middle rastergram, and the dashed line to the lower rastergram. Below the PSTHs are eye position traces, vertical position above horizontal position, with the multiple trials overlaid. At the bottom are the behavioral events. The stippled band represents the fixation point illumination, the solid band the represents the target presentation (although not apparent from this diagram, the fixation point continues to be illuminated while the target is flashed), and the sharp vertical line at $t = 0$ on the right half of the diagram represents the time of saccade. The behavioral bands consist of a set of horizontal lines, one for each trial.

Computing the principle components of the responses assumes that the underlying effects are orthogonal; we have evidence to neither support nor refute this assumption, thus we must keep it in mind when interpreting these results. This observation also lead to an attempt at improved analysis using clustering methods as described in the next section.

PSTH Clustering

By inspection, the loadings of neighboring cell PSTHs on the first four principle components exhibit no evidence of clustering when examined pairwise. This could be evidence for appropriate decomposition of orthogonal and independent axes of variation, or it could be evidence for non-orthogonality in

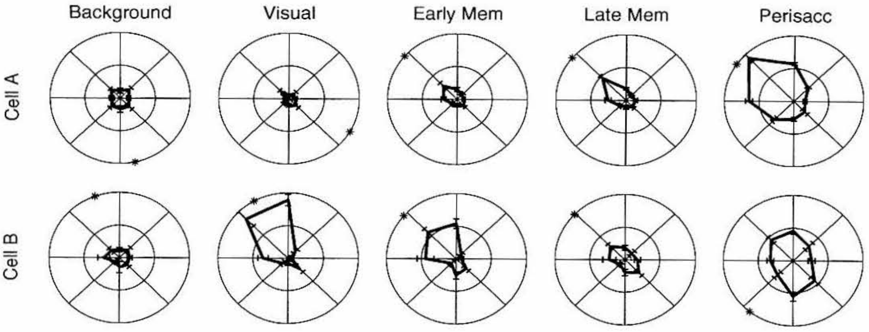

Figure 103: Two Simultaneously Recoded Cells

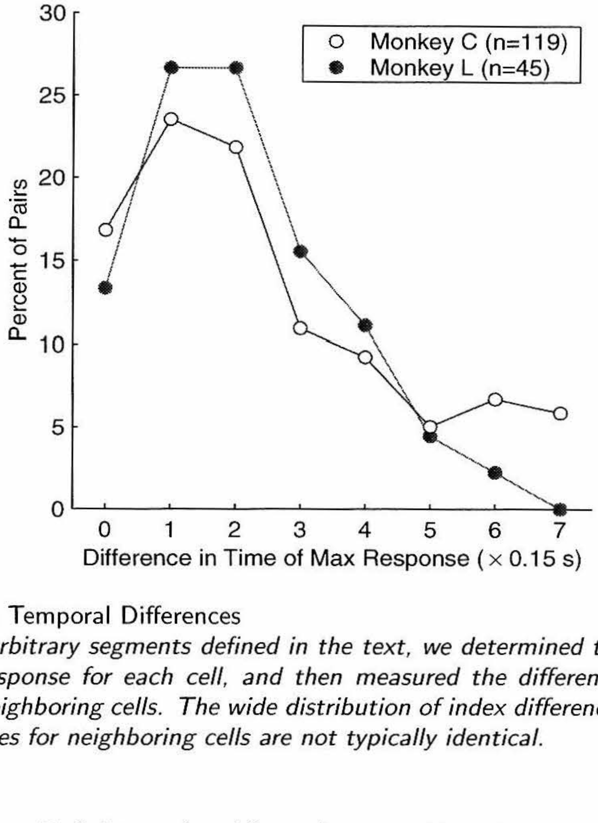

Figure 104: Temporal Differences
Using the arbitrary segments defined in the text, we determined the epoch of maximal response for each cell, and then measured the difference in epoch index for neighboring cells. The wide distribution of index differences indicates the responses for neighboring cells are not typically identical.

the underlying collection. To help resolve this, a rigorous attempt was made to cluster the PSTHs at varying binnings both with and without projecting into the principle component space using an Estimation-Maximization algorithm similar to the one used to sort spikes in our recordings. This was also intended to determine if different response classes could be deduced based on PSTH shape.

Clustering was evident neither after analysis of raw PSTH traces, nor after analysis of traces mapped to the principle component space for dimensionality reduction, although the results were numerically superior in the latter case. The attempt with the highest likelihoods is shown in Figure 106 which depicts an utter failure to detect clusters. While the figure is best viewed in color,

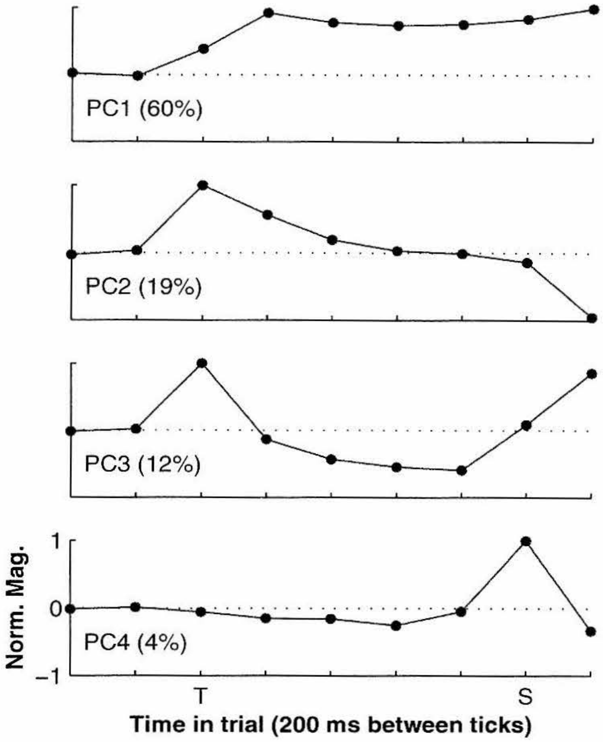

Figure 105: PSTH Principle Components
The first four principle components of the PSTH response for preferred targets. Horizontal axes are time in trial, where T marks the target flash, S marks the saccade, the memory period lies between them, and ticks indicate 200 ms increments. Vertical axes are magnitude of the component, normalized so that the maximum value of each is 1. The percent of total variance contributed by each component is shown in parentheses next to the component number.

even black-and-white reproductions should depict a mass of points without clear structure.

Although we see examples of memory responses with short and long duration, the PSTHs exhibiting memory response did not break out into separate clusters for differing lengths of sustained activity, suggesting that there is a continuum of memory response lengths.

There was an additional attempt to cluster the PSTH data using a more advanced algorithm [35] which uses a novel scaling parameter as part of the data fitting, and although not shown, the results were again equivocal. The optimal number of cluster centers appeared to be 9, but there was no clear superiority of that number over other numbers of clusters, and again, the results were sensitive to initial conditions.

3.5 Power Spectra and Oscillations

We computed the power spectrum for each of the 143 cells in the aggregate data set to look for evidence of oscillations, in particular examining the lower frequency bands (100 Hz and below). Contrary to our previous reports [53], we found evidence for oscillation in many of the cells, at

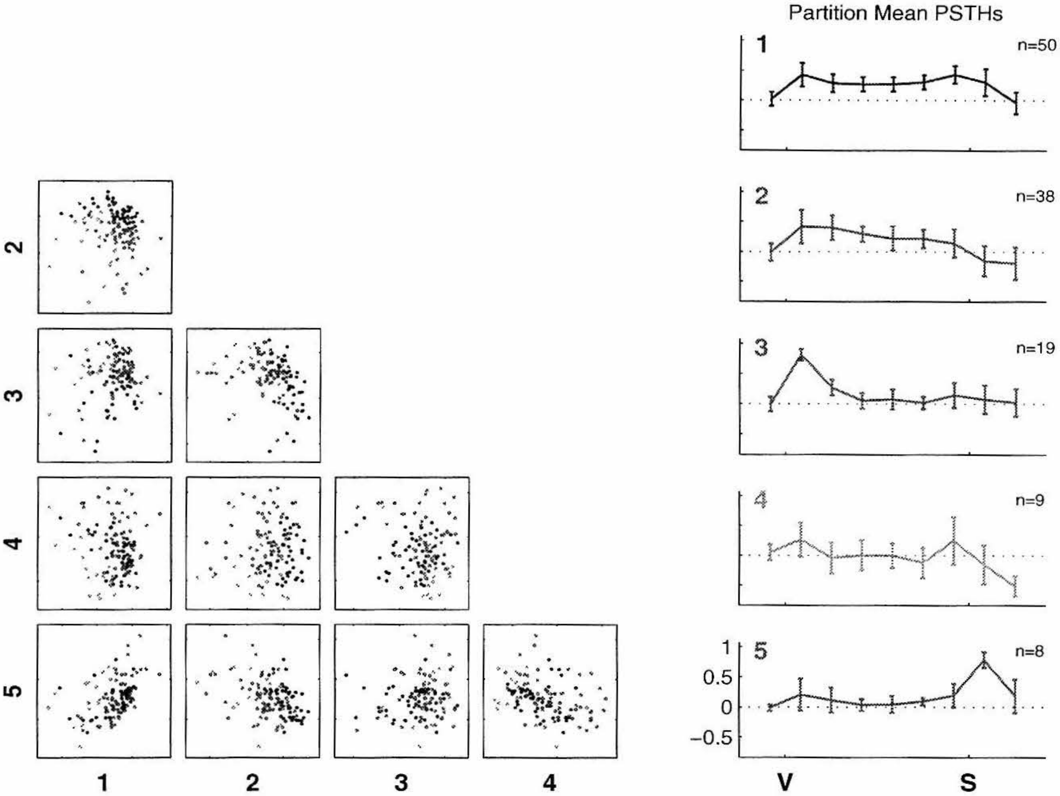

Figure 106: PSTH Partitioning
The PSTHs were projected into the principal component subspace and partitioned by an EM clustering algorithm into five groups. The panels on the left show the optimal discriminant subspace for the clustering. The clustering is soft and highly dependent upon initial conditions, thus we term it a partitioning, as there is little evidence for internal structure. The means of the elements assigned to each partition are shown on the right.

frequencies in the hundreds of Hertz, suggestive of bursting. An example spectrum from the cells in one recording is shown in Figure 107, depicting two cells with evidence for oscillation, and one without. Of the total 143 cells, 118 (82%) had spectra that significantly deviated from Poisson-simulated equivalents ($p < 0.001$, Kolmogorov-Smirnov test), and by inspection many of these, but not all, were similar to the examples in Figure 107, with regular multiple peaks suggestive of oscillatory processes. We also examined the spectra for different target conditions, and although not presented here, clear tuning effects independent of firing rate were seen.

4 Discussion

The primary difficulty with these experiments lies in the choice of using fixed-radius stimuli for the bulk of the data. Although there is substantial published evidence for similarity between nearby LIP neurons' receptive fields, there has, until this study, been no comparison of neighboring cells'

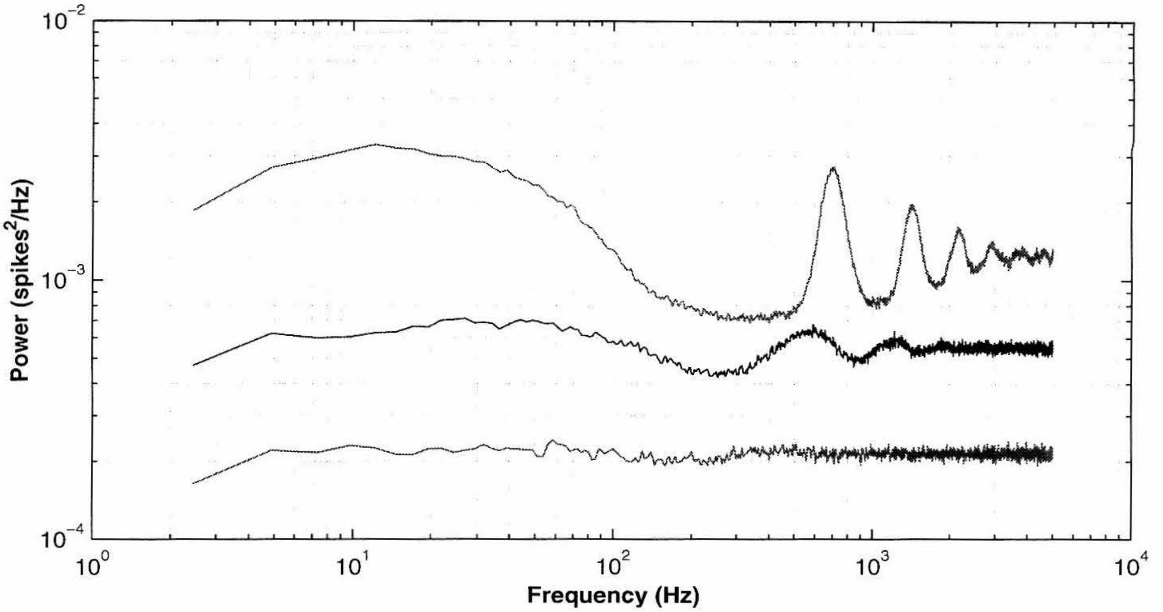

Figure 107: Example Spike Train Spectra
Spectra from three cells in one recording (lmem2306). Spike trains were binned to 0.1 ms, and segments corresponding to the period from one second before each target flash through 5 seconds after used to compute the power spectra. The spectra were computed for each trial using 4096-element overlapping windows (Hamming tapered) and averaged together across trials. The period used to compute the spectra corresponds to 89% of the total recording in this example. Repeated peaks represent the fundamental and harmonics of oscillatory processes.

receptive fields. This distinction may seem subtle except in light of results on V1 receptive fields for neighboring neurons by Maldonado and Gray [39] and Normann [46] report, which suggest a higher variation in receptive field characteristics than expected from research such as Hubel and Weisel's seminal work [29]. Given a known similarity but uncertain variation in receptive field position and size, coarse spatial sampling is perhaps a poor choice of experimental paradigm. This concern becomes even more important in the subsequent chapter when we examine fine time scale interactions between spike trains.

4.1 Preferred Direction

Previously reported figures for LIP receptive field sizes, positions, and shapes [9, 8, 59] suggest that sampling at a fixed radius such as 10° will accurately represent the direction of the receptive field peak, or, alternately put, that receptive fields are reasonably symmetric. However, it is dangerous to think the preferred direction as computed here is more meaningful than it is: a proxy for a proper measurement of the receptive field center.

We might ask why preferred directions should match for neighboring cells. We posit that there is sufficiently complex computation required within LIP that more than one neuron is required for a given output, and, therefore, we will find multiple neurons, each with slightly different dynamics of response, which will have similar, if not identical, receptive fields. By arguments on maximal leveraging of available resources, these neurons should be physically proximate. This would imply that neighboring neurons should have extremely similar receptive fields, and while we see a good match between neighboring preferred directions, it is not an exact match. Part of the discrepancy surely is from experimental uncertainty, and part from the straightforward (model-free) means we use to compute the preferred direction. Additional analysis would be necessary to determine if these errors in estimation are sufficient to explain a model with absolutely tight receptive fields.

We might also ask why are there so many receptive fields which don't match well. This has two potential answers. The first is based on the observation that the large discrepancies between preferred directions are often from cells with inhibitory or minimally excitatory tuning; these cells will have increased estimation error. Also, there is some evidence for a patchiness of topography in LIP which would lead to jumps in tuning near the edges of patches. The evenness of probability for differences in preferred direction beyond $\pi/4$ radians suggests a uniform background scatter, perhaps from experimental uncertainty, on top of which lies a narrow distribution of tightly matched cells.

4.2 Response Shapes

While we have an intuitive notion of differing and independent phases of neural response throughout the memory saccade task, there is little objective evidence in our recordings for such. While there are cells with different levels of response at each phase, there is no evidence for disparate classes of cells: those which only have visual responses, memory responses, or perisaccadic responses. While previous reports from our laboratory [8] have suggested such distinct classes, the definitions were arbitrarily imposed; we could not find compelling evidence, other than the weak implication from our principle components analysis, for such phases.

One possible confound is from not having placed our targets at the center of neural receptive fields. If the temporal profile of a neuron were merely scaled to different levels depending on the target position within its receptive field, this would not be a concern. However, if the temporal profile varies within the receptive field, especially in a continuous fashion, then this confound would lead to the inability to cluster that we have seen. The one multi-radius mapping example presented, with evidence for a response field which shrinks through the trial, supports this hypothesis, as stimulation at the center of the receptive field shows a broad, strong, and even visuo-memory response, while stimulation progressively off the center shows a memory response which weakens more quickly than the visual response. This one example does not prove the hypothesis, but suggests that additional examination of these issues is required for fuller understanding.

4.3 Visual Receptive Field Mapping

A key component therefore to future research will be the ability to quickly assess the receptive fields of a neuron or set of neurons to guide additional experimentation. While this approach to on-line tuning of an experimental paradigm is not new, nor is the idea of a reverse-correlation task even for simultaneous multiple neuronal mappings [15], the reasons for using it become more compelling with the larger numbers of neurons being monitored. Because the memory saccade paradigm requires trial times of a number of seconds, it is not possible to sample a reasonable area of visual space completely, nor would it likely be enlightening to do so without first developing a currently lacking basic understanding of the circuitry within LIP. It would be far better to accurately adjust, for example, one target at the center of each detected neuron's receptive field with one or more control targets in the contralateral hemifield. Doing this requires a swift estimation of these centers, and we have shown some evidence that this is possible. What remains is further verification of the technique.

The dynamics in the mapping task, apparently replicated in our one example in the memory saccade task, present a separate and potentially rich source of data for analysis. The latency of the second neuronal volley being beyond nominal visual latencies in the example is evidence for communication from an indirect pathway, via the frontal eye fields, and combined with simultaneous recording from that area could provide a greater understanding of the signal paths between the parietal and prefrontal areas.

5 Summary

As the first of two broad phases of analysis, we have presented results from examination of wider temporal scale responses in LIP to a memory saccade task. We presented the spatial and temporal tuning for simultaneously recorded cells and found that such neighboring neurons tended to have similarity in spatial tuning, yet dissimilarity in temporal response. These results are evidence for locally heterogeneous circuitry within parietal cortex, specifically area LIP, that carries locally homogeneous spatial selectivity.

My whole brain was out of tune.
My whole brain was out of tune.
I don't know how to tune a brain, do you?
Went into a brain shop,
They said they'd have to rebuild the whole head.
I said, "well, do what you gotta do."
— MORPHINE (My Brain, *B Sides and Otherwise*, 1997)

Chapter Eleven: Covariation and Coherence

This chapter continues analysis begun in the previous chapter on recordings made with the same memory saccade paradigm, extending the rate-based analysis to fine time-scale interactions between neighboring spike trains using tools such as covariance and coherence. This chapter is based in part on previously published work [55, 58].

1 Covariation

Action potentials from small groups of physically adjacent neurons were recorded from the parietal cortex of two rhesus macaques performing a memory saccade task. Recordings were make using tetrodes and sorted into spike trains from individual cells. Auto- and crosscovariograms of spike times for individual cells and simultaneously recorded pairs of cells, respectively, show modulations synchronized with behavioral events consistent with visual, memory, and perisaccadic activity.

1.1 Introduction

Computation within a single cortical area is likely to be a dynamic process, involving local recurrent circuitry, which can be revealed through simultaneous multiple single unit recording. Of particular interest are cells that lie within tens of microns of each other. These cells are often members of single functional unit (a column or micro-column), and are likely to have connectivity to sustain synchronous coding or cooperative computation. The lateral intraparietal area (LIP), as an example, shows a variety of responses to visual stimuli, saccade planning, and saccde execution amongst neighboring cells [55], suggesting that it contains rich local circuitry.

For this experiment, we sought to examine dynamics in spike firing for individual neurons and between neighboring cells within LIP. We used the tetrode technique as introduced by Recce and O'Keefe [60], adapted to the awake, behaving monkey preparation [54, 56] to simultaneously collect spike trains from multiple neighboring cells. We recorded from the intraparietal sulcus of two rhesus macaques while the animals performed memory saccades, and computed covariograms from the collected data to track changes in neural response.

1.2 Methods

The behavioral task for this experiment was the memory saccade [55]. Stimuli were points of light approximately 1° in diameter projected on a tangent screen, the first of which to appear for each trial was a central fixation point. The animal was required to foveate this light as long as it was illuminated. 1–2 s after fixation commenced, a second point, the target, was flashed for 150 ms in one of eight equally spaced positions around a circle of radius 10° centered on the fixation point.

After a 1000 ms memory period delay, the fixation point was extinguished. At this cue, the animal was required to saccade in the dark to the remembered target location. Upon successful completion of the cued saccade, the target was reilluminated for 500 ms, and the animal rewarded with a drop of juice. Targets were randomly interleaved and sets of 96 or 120 trials collected at each recording site.

Data were recorded from the parietal cortex from each of two animals (with two hemispheres total) using tetrodes [54]. Stereotaxic coordinates and observed neural responses were consistent with recordings having been made in the lateral intraparietal area (LIP), although recording sites have been histologically verified in only one of the two animals. Simultaneous spike trains were derived from continuous analog recordings digitized to 16 bits and sampled at 16, 20 or 24 kHz using previously-reported statistical techniques [69, 68].

Sets of auto- and crosscovariograms [11], binned to 1 ms over a delay of ±50 ms, were computed at each target position for each cell and pair of simultaneously recorded (neighboring) cells for non-overlapping 200 ms windows spanning the range of behavioral time. Each covariogram was computed by forming the correlation of the two spike trains, subtracting the shift predictor, and normalizing by an estimate of the variance given the null hypothesis of independence.

Covariogram sets were rotated by target to bring the preferred directions into registration at a fixed position (Column 3, as shown in Figure 108). For autocovariograms, the preferred direction was determined from the target with maximum first order firing rate (the spike count over time, or, equivalently, the integral of the PSTH) for the experimentally relevant epoch; for crosscovariograms, the preferred direction was determined similarly for the second order firing rate (the integral of the convolution of the two PSTHs). The central bins on all covariograms were suppressed, and mean auto- and crosscovariograms were computed over all cells and pairs of neighboring cells, respectively, after preferred-direction registration. Finally, for the preferred direction, first and second order mean firing rates were computed in the same 200 ms windows as the covariograms.

1.3 Results

The mean autocovariogram is shown in Figure 108. All of the subplots exhibit a refactory period with flanking positive side lobes 1–2 ms wide at the peak, and a few exhibit negative side lobes approximately 20 ms wide, at slightly larger delays. The inner positive side lobes range in size, but are largest for the memory period, found between visual target presentation and the saccade, in the preferred direction. Further, there is a suppression over all directions during the response to the visual target. The negative side lobes are nearly non-existent except during the memory period for the preferred direction, and for the last two or three post-saccadic windows in the opposite direction.

The mean crosscovariogram is shown in Figure 109. Most subplots show a tendency towards a broad, if shallow, central peak 30–50 ms wide. Subplots in the preferred direction covering the memory period (between **V** and **S** in the figure) show a narrow central peak, 1–2 ms wide, which

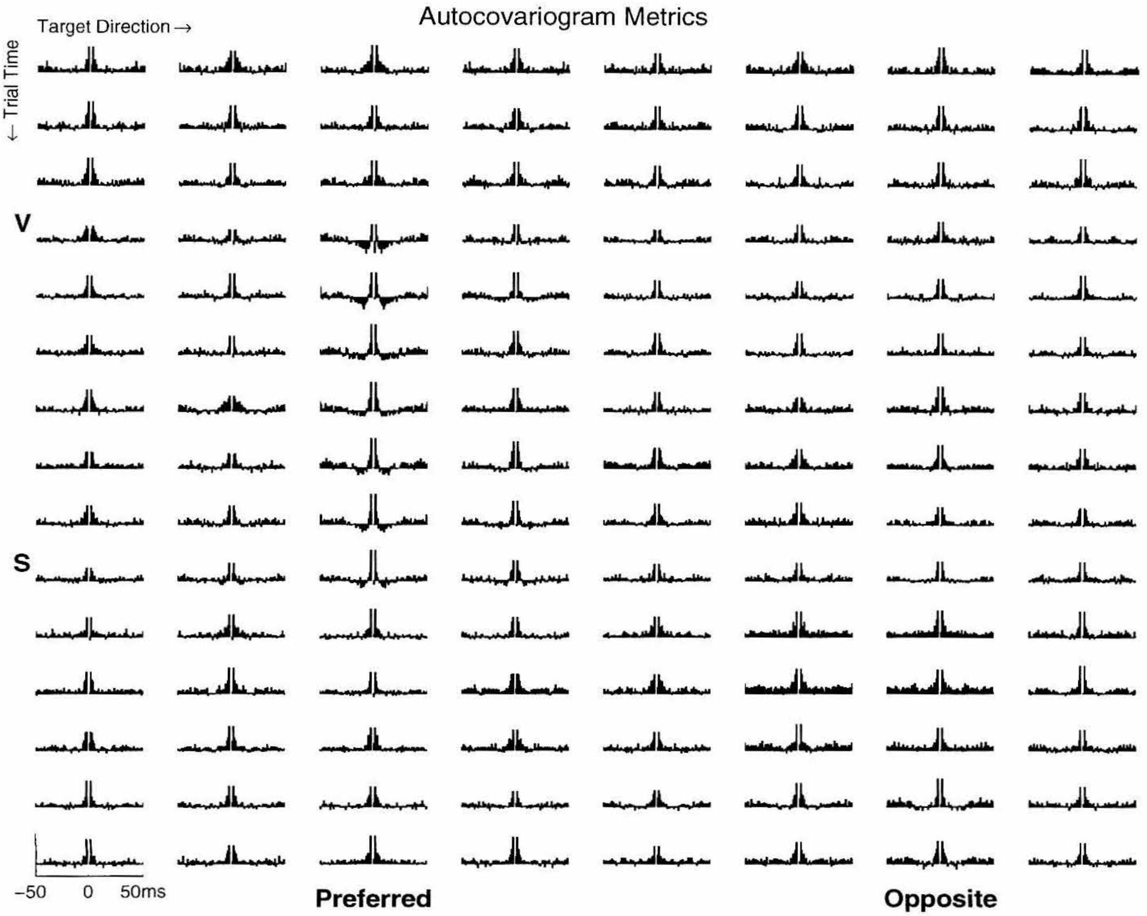

Figure 108: Mean Autocovariogram
Each column is for one target direction after registration, with the preferred direction in the third column, and target positions moving around the circle from left to right. Each row corresponds to one 200 ms window, with trial time increasing from top to bottom along rows. Row **V** *is the first window containing visual target response; row* **S** *is the first window after the saccade. From the start of the trial to the row before* **S**, *trials were temporally aligned to the target presentation, from that row onwards, trials were aligned to the measured saccadic event. Column* **Preferred** *was used to register preferred directions; column* **Opposite** *is 180° away. Each autocovariogram is binned at 1 ms and covers ±50 ms of delay. Vertical scales in arbitrary units related to significance.* n = *124 cells.*

weakens shortly after the saccade. This same peak can be observed in other subplots as well, as seen for the opposite direction towards the end of trial time.

The mean first and second order response along the preferred direction are shown in Figure 110. Both exhibit a primary peak directly after target presentation followed by sustained activity during the memory period and a secondary peak spanning execution of the saccade, while sustained memory activity is slightly stronger for the first order metric.

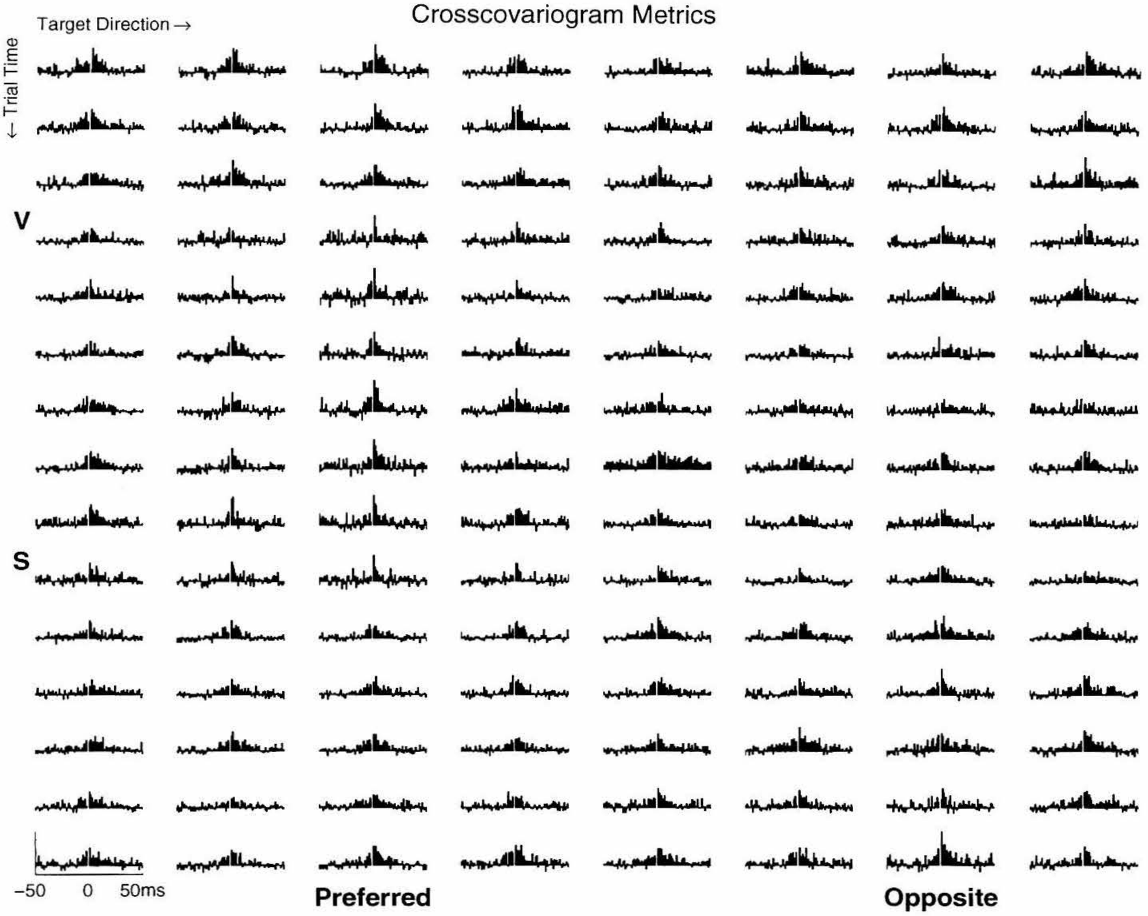

Figure 109: Mean Crosscovariogram
Columns, rows, and axes as described in the previous figure. n = 169 pairs of neighboring cells from 37 sites.

1.4 Discussion

Response-locked changes in autocovariogram profiles over time and target position are relatively straightforward to interpret. The increase in the central peaks of the autocovariograms along the preferred direction start one 200 ms window after the response to the visual target begins, and continue through to about one window after the saccade. As these peaks correspond to bursts of spikes, their growth during the preferred-direction memory period reflects a rise in the proportion of spikes occurring in bursts despite the increased overall firing rate therein. Central peak modulations in this same period correspond quite well to the time a memory structure would need to retain information about the target: initialized by a volley of visual response spikes (within the first 200 ms period), it would be reset by the saccadic event.

The changes observed in the autocovariogram negative side lobes track the mean neural response for the corresponding condition: it is highest along the preferred direction, strongest at the visual response, weakened slightly during the memory response, and re-elevated perisaccadically, but also slightly elevated late in trial time for the opposite direction. This lattermost effect we interpret to be due to uncued post-trial saccades returning to the fixation point, which, for the opposite direction, will be the same retinocentrically-expressed saccade as the one from central fixation to the preferred target. If the recorded cells are oscillating at slightly different frequencies, the mean autocovariogram would exhibit negative side lobes during periods of oscillation exactly as seen; although not shown, some individual autocovariograms support this conclusion.

Similar changes in crosscovariogram profiles are less clearly interpretable. Again, we find one feature, namely a large central peak, which tracks the experimental epoch, rising in the preferred direction directly after the target presentation, and maintaining activity until approximately 200 ms after the saccade. This only coarsely matches the time course of mean firing rates for these cells. We also see a modulation of the broader central peak across target directions, present at the start of the trial, suppressed during the 200 ms after the display of the visual target, returning during the memory period, again being suppressed directly after the saccade, and finally returning once the saccade has been completed.

1.5 Covariation Conclusions

The apparent disparity in temporal response between the two autocovariogram and two crosscovariogram effects suggests multiple independent mechanisms may be at work in LIP mediating different phases of the response. Additional work with classification of cells based on responses in different parts of the task, we hope, will help us understand these.

2 Coherence

The dynamics of cell-to-cell interactions as measured by the coherence between spike trains can be used to elucidate local circuitry and phases of operation for networks of neurons. Computing the coherence between pairs of spike trains collected from macaque area LIP during a memory saccade task shows significant coherence for frequencies below 50 Hz, and time lags less than ±10 ms. The mean coherence, as examined in both the time and frequency domains for a sliding window, undergoes changes which correspond to behavioral events, but are not necessarily linked to the mean firing rates of the recorded cells. Our results suggest that the maintenance of memory activity is done through a different mechanism than sensory response and saccade generation.

2.1 Why Coherence?

What is the justification for wanting to use coherence rather than the more-often used correlation? The diagram in Figure 110 should shed some light on this issue. Suppose the goal is to discover the unknown relation between two neural signals, such as a and b in the figure. Suppose that, as

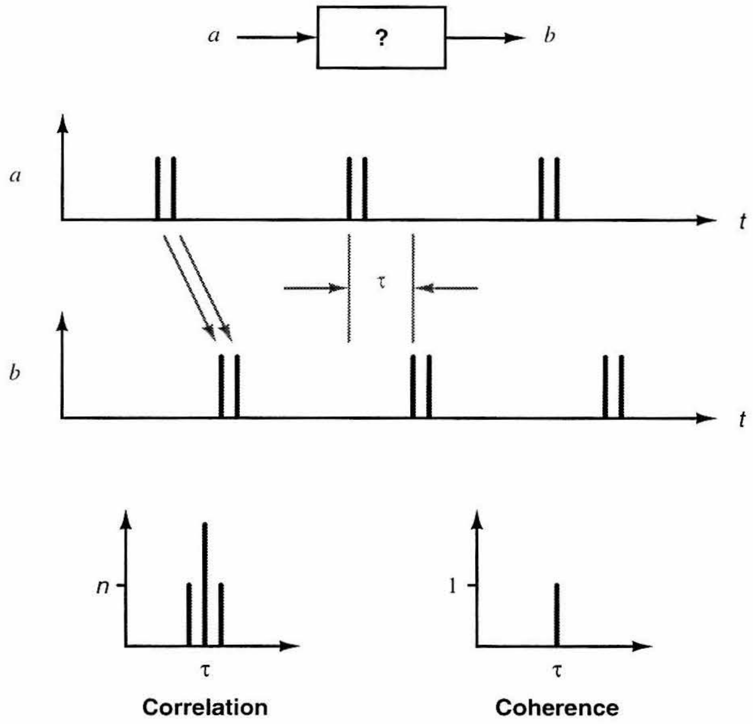

Correlation **Coherence**

Figure 110: Correlation and Coherence
The fundamental reason for using coherence as opposed to correlation can be seen in this simple example depicting the difference between the two computations on a model system which comprises a delay unit being driven by doublet spikes. The coherence between a and b shows evidence of just the delay, while the correlation confounds this with the doublet pattern in a.

omniscient beings, we already know that the relation to be that b is an exact copy of a, except delayed by period τ. Further, suppose that a consists of pairs of spikes. The correlation between a and b will show evidence of the delay τ as well as the doublet nature of a, as shown in the plot labelled Correlation,[12] while what we wish to extract is just the delay, since that is independent of any firing patterns, as shown in the plot labelled Coherence. To remove the effects of structure in a, we need to deconvolve $a \otimes b$ by the autocorrelation of a. The easiest way to do this is in the frequency domain where deconvolution becomes division; accordingly, we apply the Fourier transform to a and b, compute the correlation between them (which will be the cross-spectrum) and divide by the Fourier transform of the autocorrelations (or spectra) of a and of b. Transforming back to the time domain, we get the coherence of a and b as shown at the lower right in Figure 110. The mathematical details are presented in the next section.

---

[12] Readers familiar with mathematical correlation will realize that the full correlogram of $a \otimes b$ will extend beyond the displayed segment, with a repeated motif of decreasing amplitude in both positive and negative directions due to the finite length of a and b. This makes it more difficult to interpret the correlation, and provides a stronger argument for examining the coherence instead.

2.2 Computing Coherence

Given signals a and b from the example, which we will now explicitly express as functions of time $a(t)$ and $b(t)$, their Fourier transforms $A(\omega)$ and $B(\omega)$, and complex conjugates $A^*(\omega)$ and $B^*(\omega)$, the coherence is readily computed in the frequency domain as:

$$C_{AB}(\omega) = \frac{P_{AB}(\omega)}{\sqrt{P_{AA}(\omega)P_{BB}(\omega)}} \tag{19}$$

where $P_{AA}(\omega) = A(\omega)A^*(\omega)$ is the power spectrum of $a(t)$, and $P_{AB}(\omega) = A(\omega)B^*(\omega)$ is the cross-spectrum of $a(t)$ and $b(t)$. If $C_{AB}(\omega)$ is inverse transformed, the corresponding time-domain $c_{ab}(\tau)$ is generated.[13]

For this study, the spectra $P_{AA}(\omega)$, $P_{BB}(\omega)$, and $P_{AB}(\omega)$ were computed for mean-corrected spike trains using multi-taper estimation methods (Percival and Walden, 1993) on 512 ms windows using discrete prolate spheroidal sequences ($\text{N} = 128$, $\text{NW} = 2$) every 64 ms over the trial time span. Coherences were then computed from trial-averaged spectra for each trial condition. The formulas below detail the calculations for spike trains $a(t)$ and $b(t)$, window offset t_0, and trial condition (target position) g:

$$c_{ab}(\tau)\Big|_{t_0} = \mathcal{F}^{-1}\left[\frac{P_{AB}(\omega)}{\sqrt{P_{AA}(\omega)P_{BB}(\omega)}}\right]_{t_0} \tag{20}$$

$$P_{XY}(\omega)\Big|_{t_0} = \frac{1}{|S_g|}\sum_{s\in S_g}\left[X_s^*(\omega)Y_s(\omega)\right]_{t_0} \tag{21}$$

$$X_s(\omega)\Big|_{t_0} = \frac{1}{|W|}\sum_{w\in W}\mathcal{F}\left[w(t)\left(x_s(t-t_0) - \overline{x}(t-t_0)\right)\right] \tag{22}$$

$$\overline{x}(t) = \frac{1}{|S_g|}\sum_{s\in S_g} x_s(t) \tag{23}$$

where $\mathcal{F}$ is the Fourier transform operator, W is the set of tapers, S_g is the set of trials for target g, and X, Y, x, and $\overline{x}$, are metavariables that can be A, B, a, b, $\overline{a}$, or $\overline{b}$, as appropriate. Subtracting, for example, $\overline{a}$ from a_s corrects for the mean spike train, and will remove coherences due to low-frequency comodulation of the spike trains such as expected from mean-rate stimulus response.

2.3 Results

The magnitude of the mean coherence over all targets and all cell pairs is shown in the frequency and time domains in Figure 111. Along with the coherence for the collected spike trains, the figure also contains the results from a control computation were the same calculations were performed on simulated inhomogeneous Poisson spike trains with rates matching the recorded data. In the

---

[13] For notational clarity, we will use τ for the time-domain coherence independent variable, reserving t for use in the untransformed time-domain signals and in the guise of t_0 as a sliding window offset.

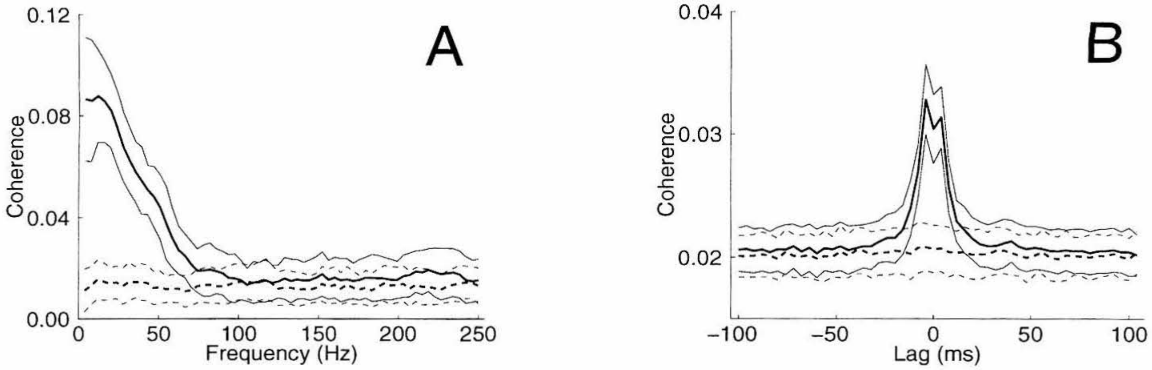

Figure 111: Mean Coherences
Frequency-domain (A), and time-domain (B) mean coherence and standard deviation (solid) of the data and inhomogeneous Poisson simulations (dashed). The vertical axis in both plots is increasing coherence. The horizontal axis in (A) is frequency, with data at approximately 2 Hz resolution, and in (B), is lag, at 4 ms resolution. These figures represent the mean coherence computed over all cell pairs ($n = 162$), all trials (10–15 trials per target), and all targets (8). The coherences are computed for each pair of simultaneously recorded spike trains on a per-target basis, then averaged over targets, averaged over pairs, and finally averaged over the experimentally-relevant epoch from target presentation to saccade completion. The frequency-domain plot shows significantly increased coherence for frequencies lower than about 50 Hz, as compared to the Poisson simulation. The time-domain plot shows significantly increased coherence for lags less than ±10 ms. The slight dip in the central time-domain peak is due to the lack of properly resolved overlapping spikes, and thus the measured value is an under-representation of the true coherence at very short lags (less than ±2 ms).

frequency domain, the primary feature is a large increase in coherence for frequencies lower than approximately 50 Hz; however, it should be noted that the mean measured coherence is everywhere larger than the mean simulated-data coherence. In the time domain, the primary feature is a central peak of increased coherence for short lags; however, it should also be noted that the mean measured coherence is everywhere larger than the mean simulated-data coherence.

Coherence over Trial Time

Mean values can hide important details, and so we begin a closer examination by expanding the plots in Figure 111 by using a sliding-window over the experimental trial, to generate the frequency domain plot in Figure 112 and the time domain plot in Figure 113. These cohereograms both display features which are locked to the behavioral events.

The mean frequency-domain cohereogram in Figure 112 exhibits features synchronized to behavioral events. Approximately 150 ms after the target flash, there is a marked decrease in low-frequency coherence which lasts 100–200 ms. Following this decrease, the coherence recovers to nearly pre-target levels for the duration of the memory period, dipping once again shortly after target reillumination.

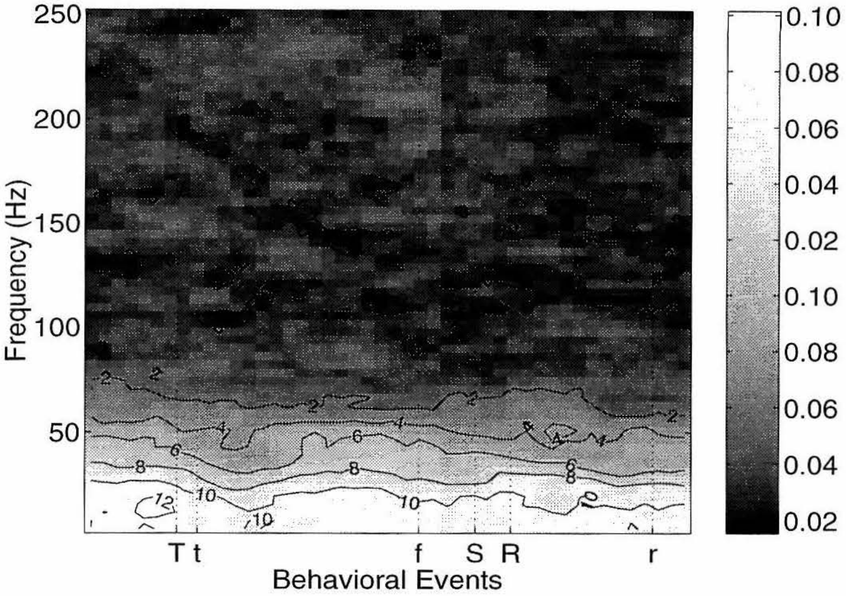

Figure 112: Frequency Domain Coherence
Mean frequency-domain coherence versus time in trial. The vertical axis is increasing frequency and the horizontal axis is time in trial. ***T*** *marks target flash onset,* ***t*** *target flash offset,* ***f*** *fixation offset,* ***S*** *the time of saccade,* ***R*** *target reillumination onset;* ***r*** *target reillumination offset, and the end of the trial. Commanded target flash duration was 100 ms; the memory period between target flash and fixation offset was 1000 ms; the whole span of the graph covers just under 3 seconds. Contour lines appear in units of standard deviation from the Poisson-simulated mean to assess point-by-point significance. Vertical dotted lines are extensions of tick marks demarking the behavioral events.*

The mean time-domain cohereogram in Figure 113 exhibits features synchronized to behavioral events like the frequency-domain cohereogram, although more subtle. There is a broad increase in coherence across all lags corresponding to the expected time-course of a purely sensory response to the target flash. Directly following this is a dip and slight narrowing in the central peak which recovers to pre-target levels by the end of the memory period. Perisaccadically there is a minor increase at all lags which lasts until shortly after the target reillumination.

Coherence over Target Condition

Important may still be hidden, so we continue closer examination by expanding the plots in figures 112 and 113 by breaking out target condition. The mean frequency-domain coherence over target condition is shown in Figure 114, which uses uncorrected spike trains, rather than PSTH-compensated data as in the previous figures. Targets were aligned according to the sum of the firing rates in each cell pair. The coherence is highest for lowest frequencies and is tuned over target condition.

The high values in the lowest bins and their modulations can be explained by PSTH-related

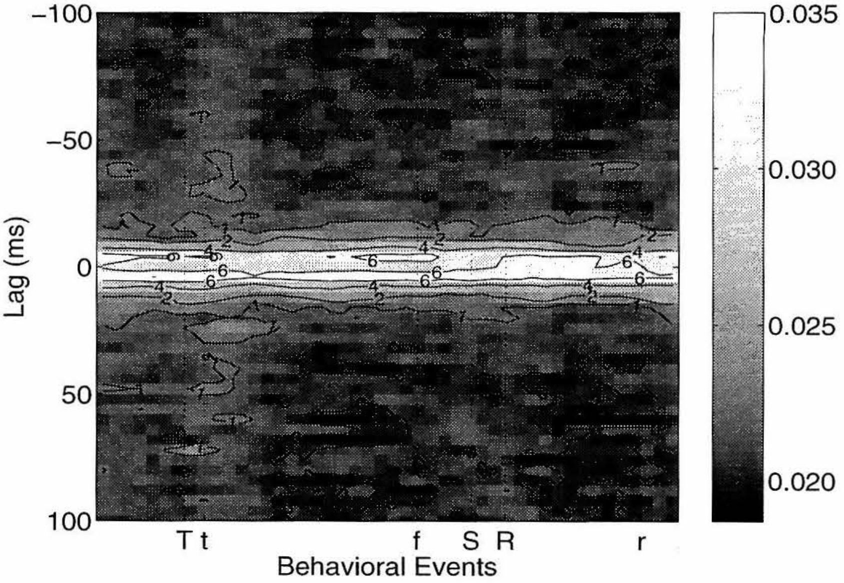

Figure 113: Time Domain Coherence
Mean time-domain coherence versus time in trial. The vertical axis is in milliseconds of lag, and the horizontal axis, contour lines, and tick extensions are as in Figure 112.

effects. This is shown for the preferred target coherence in Figure 115. We performed two control computations, The first was to calculate the coherence with trial information shuffled. For each cell pair, the trial information was permuted such that the target identity for a trial did not change, and the coherence computed otherwise as before.

The second was to calculate the coherence with varying-rate Poisson simulated spike trains. The PSTH of each cell for each trial condition was computed with 200 ms bins. A Poisson-distributed process with variable mean rate was then sampled for each trial according to condition. These synthetic spike trains were then used to compute the coherence.

Finally, we computed the coherence of spike trains after subtracting the PSTH for the appropriate trial condition from each cell. Notice that since computing coherence is a non-linear transformation, we cannot simply subtract the shuffled or Poisson-based results from the original. The two control computations retain the original's low-frequency structure, while the PSTH-corrected computation has removed it.

Examining the mean time-domain coherence over target position reveals a similar tuned process as shown in Figure 116. And, again, we performed the same computations for the preferred target response as shown in Figure 117. Here, the central peak has been eliminated in the shuffle and Poisson controls, while being retained in the PSTH-corrected computation.

Tuning effects for the time-domain coherence over target position were quantified by computing coherence for a 1 s window corresponding to the memory period first over all targets, and then

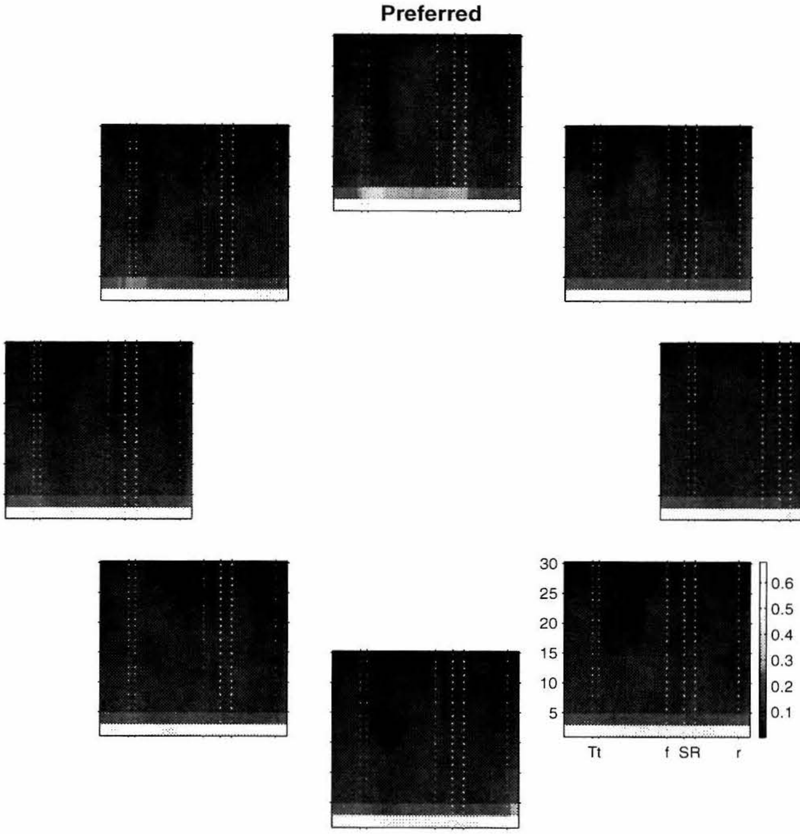

Figure 114: Coherence versus Targets (Frequency)
The population mean frequency-domain coherence, without PSTH compensation, for each target condition shows highest values at low frequencies. The horizontal axis for each subplot is time in trial, with behavioral markers as in previous figures. The vertical axis is increasing frequency, in Hertz. The low-frequency coherence is tuned for the preferred direction, but as will be seen in Figure 115, this is primarily due to comodulation of firing rate from task response rather than cell-to-cell interaction.

for each target condition. The per-target results were then normalized by the pan-target result to determine deviations from the mean response, and the result shown in Figure 118. No significant deviation over targets is seen for lags corresponding to the central peak in previous time-domain coherence figures, while significant deviation is found for the preferred direction alone for delays larger than that central feature.

2.4 Discussion

While significant coherence is seen in both time and frequency domains, it is difficult to interpret the results without additional analysis beyond the scope of this document. Keeping this caveat in mind, as well as the possibility that performing additional analysis may well modify suggested interpretations, we propose the following.

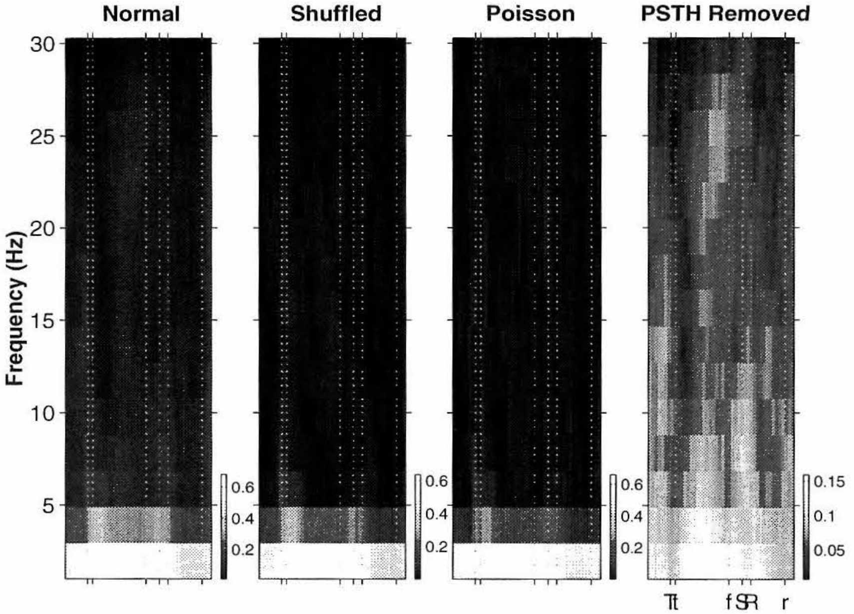

Figure 115: Low-Frequency Modulation due to PSTH Effects
The high values of low-frequency coherence are due to comodulations of firing rate, as evidenced by the two controls of the trial-shuffled coherence, and varying-rate Poisson model coherence for this data set. Once the PSTH has been subtracted from the individual spike trains, the low-frequency bins are much reduced (note differing color scales), but the higher-frequency information is retained.

Primarily, it is important to examine a coherence signal in the time domain in addition to the frequency domain. It would have been difficult to predict the peak near zero lag given only the frequency-domain coherence without performing an inverse transform. While it is true that the two domains are duals, so in a strict sense, no feature exists in one domain that is undetectable in the other, visualization in each of the two domains is a valuable tool. Indeed, the duality of time and frequency provides a ready correspondence between the coherence in one domain being everywhere larger than the Poisson control and there being a significant excursion at zero in the other domain, but the precise shape of these excursions are not easily known without actually examining them in the appropriate domain.

The observed changes in coherence suggest that different mechanisms are at work during the distinct phases of the task: pre-cue, sensory response, early memory, late memory, and peri-to-post saccade. While full understanding of the cellular mechanisms requires more detailed experimentation, these results are consistent with what would be expected from a volley of incoming sensory-response spikes that tend to desynchronize neighboring cells and trigger a memory network that takes approximately 500 ms to stabilize, in turn reset by a secondary, saccade-related volley. It should be cautioned that some of the figures presented here collapse the response over all targets,

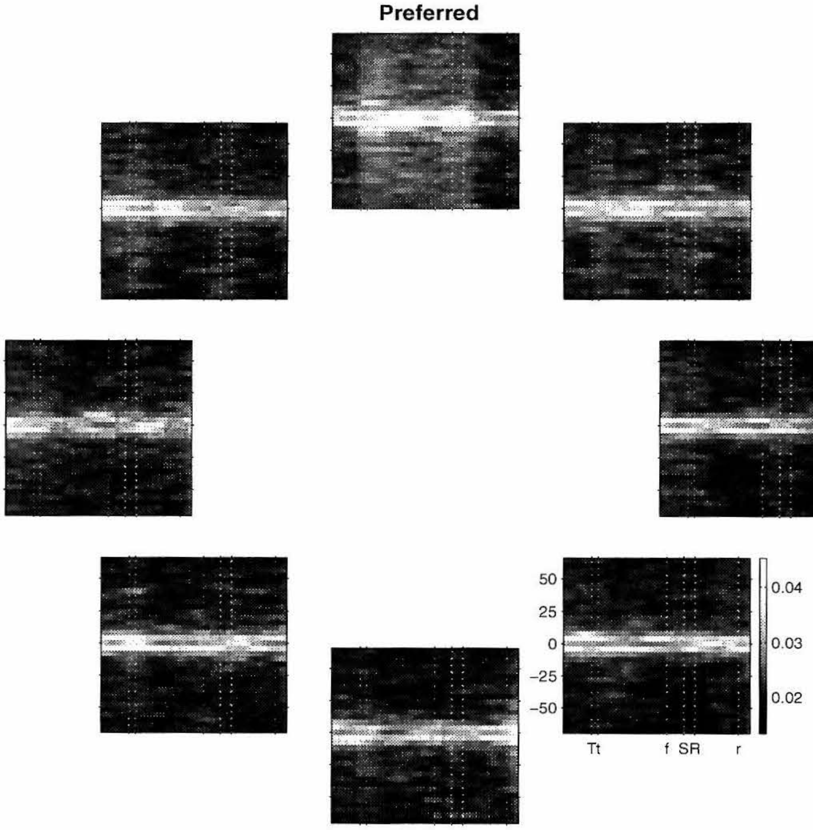

Figure 116: Coherence versus Targets (Time)
The dual of Figure 114, this shows the time-domain coherence without PSTH compensation versus target condition. The peak in central bins is evident for all target positions, while apparently stronger in the preferred target, although as discussed in the text, this modulation is not significant once the PSTH component has been removed (see also Figure 118). The coherence at flanking, delays above 25 ms is maximal for the preferred target.

and it is possible that the response for preferred targets is substantially different from non-preferred targets, requiring reinterpretation of the data.

The excess coherence during the memory period at larger delays (*ca.* 50 ms) for the preferred target, and a lack of significant tuning at shorter delays (*ca.* 10 ms) reveals changes in the underlying circuitry. These effects would be consistent with a relatively fixed interaction between neighboring cells, and a varying amount of external input to each. In particular, it indicates that for targets inside their receptive fields, neighboring cells are less likely to fire in a coherent manner at a fixed delay than for targets outside.

This would be consistent with memory circuitry that maintains activity through distributed incoherent activity. Incoherent activity would be desirable to compensate for the effects of silence during synchronized refractory periods: If all cells fire simultaneously, they will all go through a

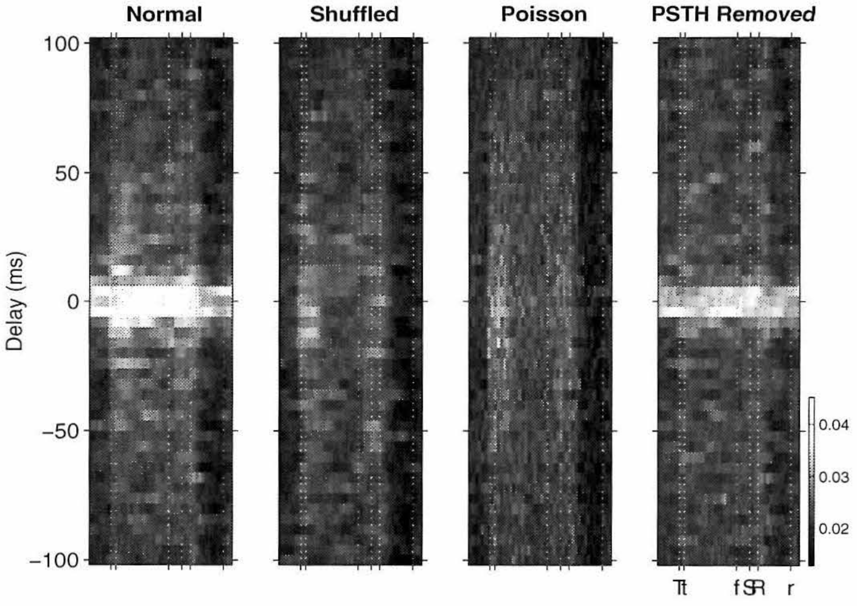

Figure 117: Shuffle and Poisson Controls
For the preferred target response shown here, shuffling trial information or using synthetic versions of spike trains retains broad scaling effects, but eliminates sharp peak in central bins. The sharp central peak is reduced, not eliminated, by removal of the PSTH.

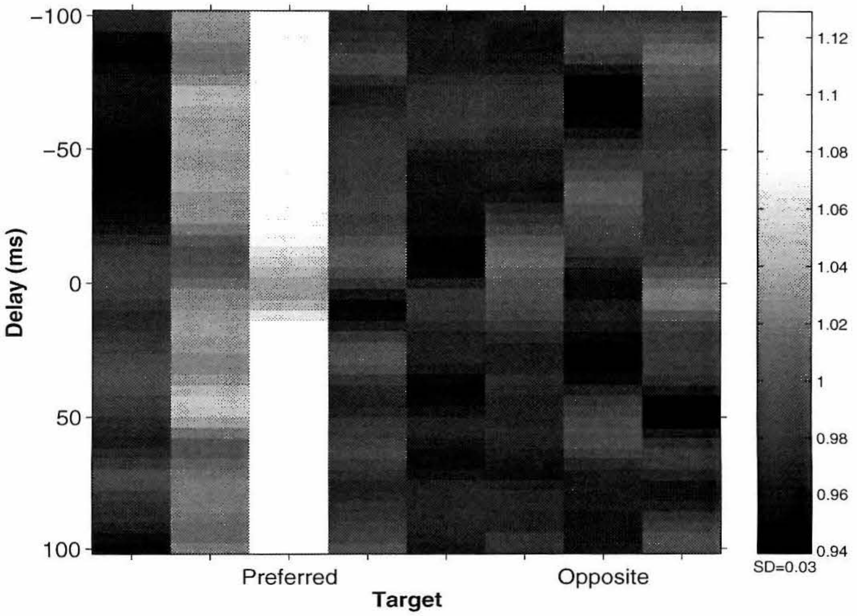

Figure 118: Normalized Memory Period Coherence versus Target
Short lag bins are untuned over target position for PSTH corrected coherences normalized to the mean coherence. Long lag bins are elevated in relation, for the preferred target.

refractory period together, at the end of which there will be no input, and the *network activity* will cease. This is consistent with the memory models proposed by Laing and colleagues [32], including the observation by Gutkin [24] that such networks can be initiated by an external excitatory input localized to a small region, and reset by an excitatory input to all cells.

2.5 Coherence Conclusions

The conclusions drawn from this section are sixfold. First, that it is useful to examine the results of coherence analysis in both time and frequency domains; using coherence instead of cross-correlation removes confounds due to input firing patterns, and produces a normalized quantity that is straightforward to compare and combine. Second, that under this memory saccade task, LIP cells exhibit significant cell-to-cell coherence at frequencies below 50 Hz. Third, that the same analysis shows significant coherence at lags less than ±10 ms. Fourth, that under this task, there are distinct phases of coherence corresponding to behavioral events: sensory, early memory, late memory, and peri-to-post saccadic. Fifth, that without the PSTH compensation used for the first two conclusions, firing rate explains most of observed coherence: Nearly all of the coherence between spike trains, when viewed both in frequency and time domains, can be explained by varying-rate Poisson spike trains. Sixth, that ±10 ms time-domain peak does not vary significantly over target condition during the memory period, however values at *ca.* 50 ms do, being significantly higher for the preferred target than other targets.

3 Summary

The two portions of this chapter, the first on correlation-based results, the second on coherence-based results, show evidence of a system within LIP which exhibits different phases of computation according to different phases of the task. The mean autocovariation function develops a secondary refractory period during the visual response, and has a higher tendency towards bursting during the memory response. The mean crosscovariation function is nearly unchanged during the visual response, save for very short latency covariation, but becomes broadly increased through the memory period. The mean coherence showed a similar effect as the crosscovariation, but most interestingly, exhibited tuning over target condition suggestive of a memory system that relies upon desynchronized firing for the maintenance of a value.

You say correlation is not causation.
— SOUL COUGHING (Moon Sammy, *Ruby Vroom*, 1994)

Chapter Twelve: Multiple Tetrodes

The distant target towards which this research was technologically aimed was recording with multiple tetrodes in the awake, behaving primate preparation with the consequent scientific ability to analyze activity from large numbers of simultaneously recorded cells. While the technical milestone was achieved, the scientific work is on-going and quantitative results are only available in preliminary form.

There are two thrusts to the continuing research, based on the number of brain areas targeted. The first uses single tetrodes in each of two brain areas, and the second uses multiple tetrodes in a single brain area. Preliminary data for the latter was presented in Chapter 4, while preliminary data for the former will be shown below.

One of the most interesting problems that has yet to be addressed is the question of scatter in receptive field location and size within LIP. While the results obtained for this dissertation indicate a strong similarity between receptive fields for neighboring neurons, a deep understanding requires finer spatial sampling than was performed. Unfortunately the memory saccade task requires a trial time of many seconds precluding the ability to densely sample visual space with multiple trials per sample point. The preliminary visual receptive field mappings shown in Chapter 10 present a starting point for this analysis, relying on previous results [8] that showed LIP visual and motor fields are aligned. The first dual-tetrode recordings used this mapping task with tetrodes in each of two hemispheres. The receptive fields for example simultaneously recorded cells from each hemisphere are shown in Figure 119, with clear contralateral tuning.

A secondary, but equally important problem is to understand interactions between LIP and the frontal eye field (FEF). Strong reciprocal and topographic projections exist between the two areas [9] leading to questions about the relative roles played in target selection, memory maintenance, and saccade generation. By knowing the topographic organization of FEF, an appropriate prefrontal location can be selected once the receptive fields of an LIP site have been characterized, such as with the mapping task detailed in Chapter 9, allowing comparisons between portions of the two areas with similar and dissimilar spatial tuning centers.

Given the proven reliability of the tetrode recording technique in macaque cortex, a broad expanse of previously unanswerable questions, ones which require multiple simultaneously recorded spike trains, can now be addressed. We look forward to the results.

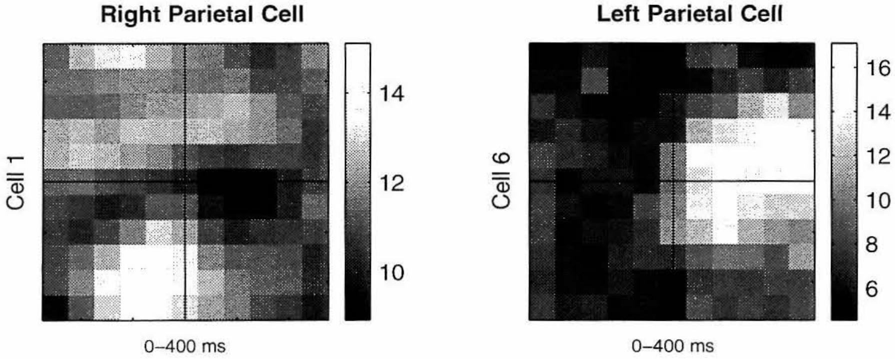

Figure 119: Simultaneously Recorded Dual Hemisphere Receptive Fields
An example from the first dual-hemisphere tetrode recordings showing receptive fields in contralateral visual hemifields for two cells, one from each hemisphere, made in an animal (Monkey Y) with bilateral parietal chambers. The left panel is the visual response spatial map for an example cell from the right parietal chamber, the right panel for an example cell from the left chamber. Firing rate is coded according to the scale bars to the right of each panel. (ydmap0304a,b)

now due to a construct in my mind
that makes their falling and their flight symbolic of my entire existence
it becomes important for me to get up and see
their last-second curves towards flight
it's almost as if my life will fall unless I see their ascent
— CAKE (Mr. Mastadon Farm, *Motorcade of Generosity*, 1994)

Bibliography

[1] M. Ahissar, E. Ahissar, H. Bergman, and E. Vaadia, *Encoding of sound-source location and movement-activity of single neurons and interactions between adjacent neurons in the monkey auditory-cortex*, J. Neurophysiology, 67 (1992), pp. 203–215.

[2] R. A. Andersen, C. Asanuma, G. K. Essick, and R. M. Siegel, *Corticocortical connections of anatomically and physiologically defined subdivisions within the inferior parietal lobule*, J. Comparative Neurology, 296 (1990), pp. 65–113.

[3] R. A. Andersen, R. M. Bracewell, S. Barash, J. W. Gnadt, and L. Fogassi, *Eye position effects on visual memory and saccade-related activity in areas LIP and 7a of macaque*, J. Neuroscience, 10 (1990), pp. 1176–1196.

[4] R. A. Andersen, G. K. Essick, and R. M. Siegel, *Neurons of area 7a activated by both visual stimuli and oculomotor behavior*, Exp. Brain Res., 67 (1987), pp. 316–322.

[5] J. S. Baizer, L. G. Ungerleider, and R. Desimone, *Organization of visual inputs to the inferior temporal and posterior parietal cortex in macaques*, J. Neuroscience, 11 (1991), pp. 168–190.

[6] R. Balint, *Seelenlähmung des 'Schauens', optische Ataxie, räumlich Störung der Aumferkasmkeit*, Monatsschr. Psychiatr. Neurol., 25 (1909), pp. 51–81.

[7] S. Barash, R. M. Bracewell, L. Fogassi, J. W. Gnadt, and R. A. Andersen, *Saccade-related activity in the lateral intraparietal area I. Temporal properties — comparison with area 7a*, J. Neurophysiology, 66 (1991), pp. 1095–1108.

[8] S. Barash, R. M. Bracewell, L. Fogassi, J. W. Gnadt, and R. A. Andersen, *Saccade-related activity in the lateral intraparietal area II. Spatial properties*, J. Neurophysiology, 66 (1991), pp. 1109–1124.

[9] G. J. Blatt, R. A. Andersen, and G. R. Stoner, *Visual receptive field organization and cortico-cortical connections of the lateral intraparietal area (area LIP) in the macaque*, J. Comparative Neurology, 299 (1990), pp. 421–445.

[10] T. Bonhoeffer and A. Grinvald, *Iso-orientation domains in cat visual-cortex are arranged in pinwheel-like patterns*, Nature, 353 (1991), pp. 429–431.

[11] C. D. Brody, *Disambiguating different covariation types*, Neural Computation, 11 (1999), pp. 1527–1535.

[12] G. Buzsáki and A. Kandel, *Somadendritic backpropagation of action potentials in cortical pyramidal cells of the awake rat*, J. Neurophysiology, 79 (1998), pp. 1587–1591.

[13] R. W. Cox and J. S. Hyde, *Software tools for analysis and visualization of fMRI data*, NMR Biomed., 10 (1997), pp. 171–178.

[14] R. C. deCharms, D. T. Blake, and M. M. Merzenich, *A multielectrode implant device for the cerebral cortex*, J. Neuroscience Methods, 93 (1999), pp. 27–35.

[15] R. Eckhorn, F. Krause, and J. I. Nelson, *The rf-cinematogram — a cross-correlation technique for mapping several visual receptive-fields at once*, Biol. Cybernetics, 69 (1993), pp. 37–55.

[16] W. E. Faller and M. W. Luttges, *Spatiotemporal analysis of simultaneous single-unit activity in the dragonfly: I. Cellular activity patterns*, Biological Cybernetics, 65 (1991), pp. 381–389.

[17] M. S. Fee, P. P. Mitra, and D. Kleinfeld, *Automatic sorting of multiple-unit neuronal signals in the presence of anisotropic and non-Gaussian variability*, J. Neuroscience Methods, 69 (1996), pp. 175–188.

[18] D. J. FELLEMAN AND D. C. VAN ESSEN, *Distributed hierarchical processing in the primate cerebral cortex*, Cereb. Cortex, 1 (1991), pp. 1–47.

[19] S. H. FUNG, D. BURSTEIN, AND R. T. BORN, *In vivo microelectrode track reconstruction using magnetic resonance imaging*, J. Neuroscience Methods, 80 (1998), pp. 215–24.

[20] P. W. GLIMCHER, M. L. PLATT, V. M. CIARAMITARO, A. HANDEL, M. A. BROWN, AND H. M. BAYER, *Real-time ultrasound imaging of neuroanatomical structure and single-unit electrodes in the living primate brain*, in Society for Neuroscience Abstracts, 1999, p. 618.12.

[21] J. W. GNADT AND R. A. ANDERSEN, *Memory related motor planning activity in posterior parietal cortex of macaque*, Exp. Brain Res., 70 (1988), pp. 216–220.

[22] P. M. GOCHIN, E. K. MILLER, C. G. GROSS, AND G. L. GERSTEIN, *Functional interactions among neurons in inferior temporal cortex of the awake macaque*, Exp. Brain Res., 84 (1991), pp. 505–516.

[23] C. M. GRAY, P. E. MALDONADO, M. WILSON, AND B. MCNAUGHTON, *Tetrodes markedly improve the reliability and yield of multiple single-unit isolation from multiunit recordings in cat striate cortex*, J. Neuroscience Methods, 63 (1995), pp. 43–54.

[24] B. S. GUTKIN, C. R. LAING, C. L. COLBY, C. C. CHOW, AND G. B. ERMENTROUT, *Turning on and off with excitation*, Neurocomputing, ((in press)).

[25] N. G. HATSOPOULOS, C. L. OJAKANGAS, L. PANINSKI, AND J. P. DONOGHUE, *Information about movement direction obtained from synchronous activity of motor cortical neurons*, Proceedings National Academy of Science, 95 (1998), pp. 15706–15711.

[26] H. HECAEN AND A. D. AJURIAGUERRA, *Balint's syndrome (psychic paralysis of gaze) and its minor forms*, Brain, 77 (1954), pp. 373–400.

[27] O. HIKOSAKA AND R. H. WURTZ, *Visual and oculomotor functions of monkey substantia nigra pars reticulara. III. Memory-contingent visual and saccade responses.*, J. Neurophysiology, 49 (1983), pp. 1268–1284.

[28] U. G. HOFMANN, D. T. KEWLEY, M. D. HILLS, S. D. V. HOOSER, G. T. A. KOVACS, AND J. M. BOWER, *Depth profiling of neural activity by multisite microelectrode arrays*, European J. Neuroscience, 10 (1998), p. 19406.

[29] D. H. HUBEL AND T. N. WIESEL, *Receptive fields, binocular interaction and functional architecture in the cat's visual cortex*, J. Physiology (London), 160 (1962), pp. 106–154.

[30] S. J. JUDGE, B. J. RICHMOND, AND F. C. CHU, *Implantation of magnetic search coils for measurement of eye position: an improved method*, Vision Research, 20 (1980), pp. 535–538.

[31] D. T. KEWLEY, M. D. HILLS, D. A. BORKHOLDER, I. E. OPRIS, N. I. MALUF, C. W. STORMENT, J. M. BOWER, AND G. T. A. KOVACS, *Plasma-etched neural probes*, Sensors and Actuators A-Physical, 58 (1997), pp. 27–35.

[32] C. R. LAING, C. C. CHOW, AND G. B. ERMENTROUT, *Stationary bumps in networks of spiking neurons*, Neurocomputing, ((in press)).

[33] D. LEE, N. L. PORT, W. KRUSE, AND A. P. GEORGOPOULOS, *Variability and correlated noise in the discharge of neurons in motor and parietal areas of the primate cortex*, J. Neuroscience, 18 (1998), pp. 1161–1170.

[34] C. R. LEGÉNDY, M. SALCMAN, AND N. BRENNAN, *A multiple floating microelectrode for chronic implantation and long-term single unit recording in the cat*, Electroencephalography and Clinical Neurophsiology, 58 (1984), pp. 285–288.

[35] J. LINDEN, *Responses to auditory stimuli in macaque lateral intraparietal area*, PhD thesis, California Institute of Technology, Pasadena, California, 1999.

[36] J. C. LYNCH, *The role of the parieto-occipital association cortex in oculomotor control*, Exp. Brain Res., 41 (1980), p. A32.

[37] J. C. LYNCH, A. M. GRABIEL, AND L. J. LOBECK, *The dirrerential projection of two cytoarchitectonic subregions of the inferior parietal lobule of macaque upon the deep layers of the superior colliculus*, J. Comparative Neurology, 253 (1985), pp. 241–254.

[38] J. C. LYNCH AND J. W. MCLAREN, *Deficits of visual-attention and saccadic eye-movements after lesions of parietooccipital cortex in monkeys*, J. Neurophysiology, 61 (1989), pp. 74–90.

[39] P. E. MALDONADO AND C. M. GRAY, *Heterogeneity in local distributions of orientation-selective neurons in the cat primary visual cortex*, Visual Neuroscience, 13 (1996), pp. 509–516.

[40] P. MAZZONI, R. M. BRACEWELL, S. BARASH, AND R. A. ANDERSEN, *Motor intention activity in the macaque's lateral intraparietal area I. Dissociation of motor plan from sensory memory*, J. Neurophysiology, 76 (1996).

[41] B. L. MCNAUGHTON, J. O'KEEFE, AND C. A. BARNES, *The stereotrode: A new technique for simultaneous isolation of several single units in the central nervous system from multiple unit records*, J. Neuroscience Methods, 8 (1983), pp. 391–397.

[42] J. MILLAR AND G. V. WILLIAMS, *Ultra-low noise silver-plated carbon fibre microelectrodes*, J. Neuroscience Methods, 25 (1988), pp. 59–62.

[43] V. B. MOUNTCASTLE, J. C. LYNCH, A. GEORGOPOULOS, H. SAKATA, AND C. ACUNA, *Posterior parietal association cortex of the monkey: command functions for operations within extrapersonal space*, J. Neurophysiology, 38 (1975), pp. 871–908.

[44] V. B. MOUNTCASTLE, H. J. REIBOECK, G. F. POGGIO, AND M. A. STEINMETZ, *Adaptation of the Reitboeck method of multiuple microelectrode recording to the neocortex of the waking monkey*, J. Neuroscience Methods, 36 (1991), pp. 77–84.

[45] A. M. NICHOLS, T. W. RUFFNER, M. A. SOMMER, AND R. H. WURTZ, *A screw microdrive for adjustable chronic unit recording in monkeys*, J. Neuroscience Methods, 81 (1998), pp. 185–188.

[46] R. A. NORMANN, *V1 receptive fields measured with a Utah array.* To be published in a Vision Research article.

[47] L. A. ORR AND E. M. LIEBERMAN, *Studies of the cardiac-like action-potential in crayfish giant-axons induced by platinized tungsten metal-electrodes*, J. Exp. Biology, 128 (1987), pp. 1–17.

[48] B. PESARAN, J. S. PEZARIS, M. SAHANI, P. MITRA, AND R. A. ANDERSEN, *Temporal structure in neuronal activity during working memory in macaque parietal cortex.* Submitted to J. Neuroscience, 2000.

[49] J. S. PEZARIS, *A reward juicer*, Tech. Rep. TR-99-01, Computation and Neural Systems, California Institute of Technology, Pasadena, California, 1999.

[50] ——, *Recording chamber drawings*, Tech. Rep. TR-00-02, Computation and Neural Systems, California Institute of Technology, Pasadena, California, 2000.

[51] J. S. PEZARIS AND D. J. DUBOWITZ, *MRI localization of extracellular electrodes using metallic deposition at 1.5T*, in International Society for Magnetic Resonance in Medicine, Philadelphia, PA., May 22 1999.

[52] J. S. PEZARIS, M. SAHANI, AND R. A. ANDERSEN, *Multiple single unit recording using tetrodes in macaque visual cortex: electrode design and spike identification*, in Society for Neuroscience Abstracts, vol. 21(2), 1995, p. 905.

[53] ——, *Extracellular recording from adjacent neurons: II. correlations in macaque parietal cortex*, in Society for Neuroscience Abstracts, 1997.

[54] ——, *Tetrodes for monkeys*, in Computational Neuroscience, J. M. Bower, ed., Plenum Press, New York, 1997.

[55] ——, *Extracellular recording from multiple neighboring cells: response properties in parietal cortex*, in Computational Neuroscience, J. M. Bower, ed., Elsevier Press, 1998.

[56] ——, *New and improved tetrodes for monkeys*, Tech. Rep. TR-98-03, Computation and Neural Systems, California Institute of Technology, Pasadena, California, 1998.

[57] ——, *A chronic chamber microdrive for neural recording.* application to the US Patent Office (CIT 2861), March 5 1999.

[58] ——, *Response-locked changes in auto- and cross-covariations in parietal cortex*, in Neurocomputing, vol. 26–27, 1999, pp. 471–476.

[59] M. L. PLATT AND P. W. GLIMCHER, *Reponse fields of intraparietal neurons quantified with multiple saccadic targets*, Exp. Brain Res., 121 (1998), pp. 65–75.

[60] M. L. RECCE AND J. O'KEEFE, *The tetrode: an improved technique for multi-unit extracellular recording*, in Society for Neuroscience Abstracts, vol. 15(2), 1989, p. 1250.

[61] R. C. REID AND J. M. ALONSO, *Specificity of monosynaptic connections from thalamus to visual-cortex*, Nature, 378 (1995), pp. 281–284.

[62] H. J. REITBOECK, *Fiber microelectrodes for electrophysiological recordings*, J. Neuroscience Methods, 8 (1983), pp. 249–262.

[63] D. A. ROBINSON, *A method for measuring eye movement using a scleral search coil in a magnetic field*, IEEE Trans. Biomed. Eng., 10 (1963), pp. 137–145.

[64] D. A. ROBINSON, *The electrical properties of metal microelectrodes*, Proceedings of the IEEE, 56 (1968), pp. 1065–1071.

[65] D. L. ROBINSON, M. E. GOLDBERG, AND G. B. STANTON, *Parietal association cortex in the primate: sensory mechanisms and behavioral modulation*, J. Neurophysiology, 41 (1978), pp. 910–932.

[66] P. J. ROUSCHE, R. S. PETERSEN, S. BATTISTON, S. GIANNOTTA, AND M. E. DIAMOND, *Examination of the spatial and temporal distribution of sensory cortical activity using a 100-electrode array*, J. Neuroscience Methods, 90 (1999), pp. 57–66.

[67] M. SAHANI, *Latent Variable Models for Neural Data Analysis*, PhD thesis, California Institute of Technology, Pasadena, California, 1999.

[68] M. SAHANI, J. S. PEZARIS, AND R. A. ANDERSEN, *Extracellular recording from multiple neighboring cells: A maximum-likelihood solution to the spike-separation problem*, in Computational Neuroscience, J. M. Bower, ed., Plenum Press, New York, 1998.

[69] ——, *On the separation of signals from neighboring cells in tetrode recordings*, in Advances in Neural Information Processing Systems 10, M. I. Jordan, M. J. Kearns, and S. A. Solla, eds., MIT Press, Cambridge, MA, 1998.

[70] D. W. TANK AND D. KLEINFELD, *The relation between electrode placement and the amplitude of extracellularly-recorded action-potentials*, Biophysical Journal, 49 (1986). Part 2.

[71] S. VENKATACHALAM, M. S. FEE, AND D. KLEINFELD, *Ultra-miniature headstage with 6-channel drive and vacuum-assisted micro-wire implantation for chronic recording from the neocortex*, J. Neuroscience Methods, 90 (1999), pp. 37–46.

[72] M. WEHR, J. S. PEZARIS, AND M. SAHANI, *Simultaneous paired intracellular and tetrode recordings for evaluating the performance of spike sorting algorithms*, in Neurocomputing, vol. 26–27, 1999, pp. 1061–1068.

[73] M. A. Wilson and B. L. McNaughton, *Dynamics of the hippocampal ensemble code for space*, Science, 261 (1993), pp. 1055–1058.

[74] M. A. Wilson, B. L. McNaughton, and K. Stengel, *A high-density miniature microdrive array for chronic stimulation and recording of patterned multiple single unit activity in the rat hippocampus*, in Society for Neuroscience Abstracts, vol. 17, 1991, p. 1395.

Oh, I'm tired, I'm tired,
I've had my fill
— TRIBE (Joyride, *Abort*, 1991)

My head, my head
Did you see my head roll away?
— HERETIX (My Head, *Gods & Gangsters*, 1990)

I think my brain wants me dead.
— POOKA STEW (Out of My Head, *In Our Minds*, 1996)