Cognitive Neural Prosthetics: Brain Machine Interfaces Based in Parietal Cortex

Thesis by

Daniella Meeker

In partial fulfillment of the requirements for the degree of Doctor of Philosophy

California Institute of Technology
Pasadena, California

Defended February 21, 2005

Acknowledgements

I thank my advisor, Richard Andersen, for his support, direction, and patience since I joined his lab. The members of this lab have been generous with their knowledge and time; they have helped me immensely over the years, and I feel incredibly fortunate to have studied in such a rich intellectual climate comprised of so many brilliant individuals. In particular, I thank Krishna Shenoy, who by way of his mentorship, example, and collaboration during my first years at Caltech has continued to inspire me as a model for scientific and academic excellence. I thank Shiyan Cao, who also collaborated on the work in Chapters 2 and 3 and who spent many wee hours working with me developing hardware and software. I am grateful to Boris Breznen, who has shared equipment, knowledge, and many encouraging words with me over the years. Additional data supporting Chapters 5 and 6 were kindly supplied by Sam Musallam. I also thank Betty Gilliken, Kelsie Pejsa, David Dubowitz, Sohaib Kureshi, Igor Fineman, and Hans Scherberger for expertise in animal care, imaging, and surgery, particularly in several early attempts to implement electrode array surgeries.

I thank the members of my committee, Shin Shimojo, Mark Konishi, and Pietro Perona for their aid and time, and especially Joel Burdick for his extensive support during every phase of this work.

I thank Cierina Reys and Tessa Yao for their administrative assistance and Viktor Schherbatyuk for computer support.

I would finally like to thank my friends and family for the many direct and indirect ways in which they have helped and inspired me over time.

Table of Contents

List of Figures

1 Introduction

The term "cybernetics" was coined by Norbert Weiner in 1947 in his book *Cybernetics: or, Control and Communication in the Animal and the Machine*. His foresight in this work predicted the ever closer interaction between the products of electrical engineering and biology. In the 1970s the notion of fusion between man and machine was popularized in the TV series *The Six Million Dollar Man*, where Col. Steve Austin is saved from certain paraplegia by bioengineering, becoming the first [albeit fictional] bionic man. Since then there has been a continuing fascination with the notion of physical integration between man and machine. The recent surge of progress in the field of Brain Machine Interfaces has brought what was once considered a castle in the sky down to a reality that may be achieved in our lifetimes.

Though there are many limitations to overcome, the study of the brain from an engineering perspective, that is, with the intent of applying the neural information rather than merely characterizing it, provides for new and better understandings of both the biological and technical problems at hand. We not only get better ideas of the accuracy and limitations of neural information, but are given a means of understanding dynamics that may not have been observable previously, such as how the brain changes with the use of BMIs. Chapter 2 sets up the BMI problem and demonstrates a proof of concept of a BMI using data from a former experiment showing that spikes in the Parietal Reach Region (PRR) are tuned to planned reach direction. Chapter 3 documents the instantiation of such a BMI and its effect on neural coding. Chapter 4 shows that non-spiking neural data, the Local Field Potential (LFP), can also be used for BMI control. The results of the experiments outlined in Chapters 3 and 4 give rise to questions about the nature of learning and neural adaptation in the Parietal Cortex. Chapter 5 looks at some preliminary results on the relationship between motivation, timing of spikes and LFPs, and spike timing dependent plasticity (STDP), and presents a model to inspire further experiments. Chapter 6 briefly describes the importance of these experiments in the context of neural prosthetics and neuroscience in general.

1.1 Motivation for PRR Based BMIs

1.1.1 The Dorsal Stream

An important distinction has been made between two streams of visual processing in the brain, delegating to the ventral stream object recognition and the dorsal stream spatial representation (Ungerleider and Mishkin 1982). The dorsal stream was further characterized as specializing its use of spatial information for action, such as navigation (MT/MST), or visual and/or reaching target selection by Goodale and Milner (Goodale and Milner 1992). The flow of the dorsal stream from the primary visual cortex to its endpoint in the parietal cortex transfers visual information first to V5 and MT/MST (optic flow (Maunsell and Newsome 1987)) and V6, a.k.a. parieto-occipital area (PO) via V2 (Colby, Gattass et al. 1988). The information is then conveyed to MIP (which houses PRR) (Blatt, Andersen et al. 1990), 7a (eye-movements (Mountcastle, Lynch et al. 1975)), AIP (grasping (Murata, Gallese et al. 1996)), and areas LIP (eye-movements, (Gnadt and Andersen 1988)) and 7b (hand-movements (Hyvarinen and Poranen 1974)). V6 has been further subdivided into V6A and V6, which share cytoarchitectonic and functional properties most similar to parietal cortex and visual cortex, respectively (Galletti, Fattori et al. 1997; Rizzolatti 2003). The segregation between eye and hand movements is maintained downstream of the parietal cortex in prefrontal areas, indicating an effector-based segregation of visual information.

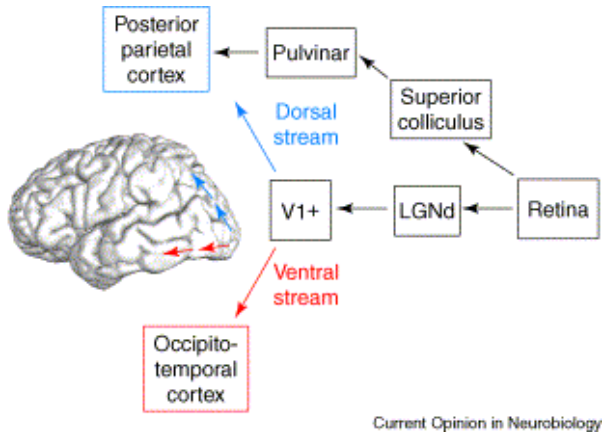

Figure 1-1 Division of the dorsal and ventral streams. From (Goodale and Westwood 2004). There are two alternate routes from the retina to PPC, one via the Superior colliculus and Pulvinar, the other via the Dorsal stream of visual processing from V1. The Ventral stream of visual information conveys information from V1 to the occipito-temporal cortex.

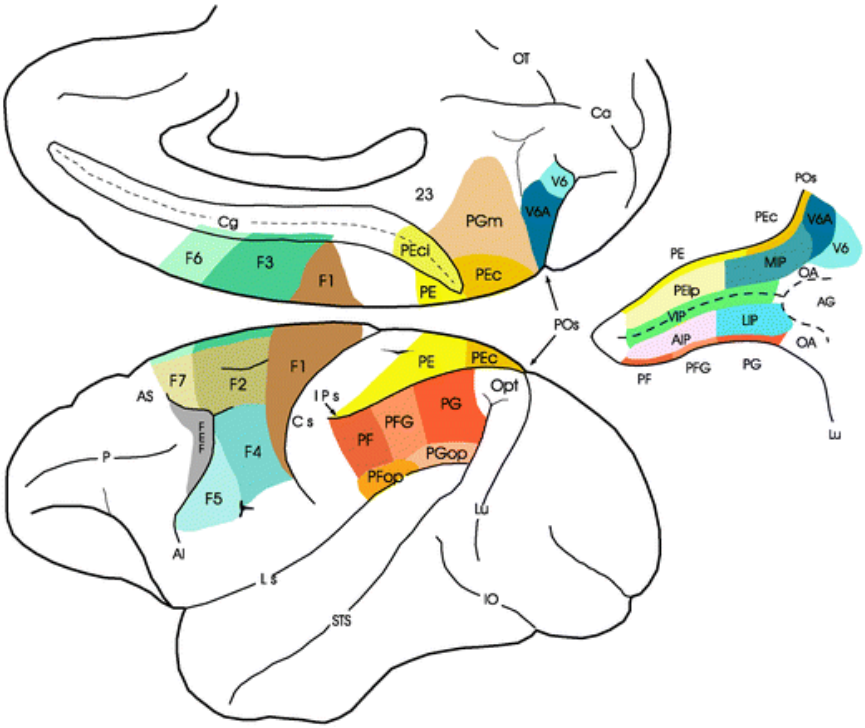

Figure 1-2 From (Rizzolatti 2003). View of elements of the dorsal stream and its targets. The IPS, where MIP/PRR and LIP are located, is unfolded to the right of the lateral depiction of the brain.

It is of note that observations of segregation between arm and eye movements in areas like the MIP (Batista, Buneo et al. 1999) and the LIP (Toth and Assad 2002) are dependent on a distinction between cues for such targets. Presumably the *recognition* (i.e. the assignment of meaning to visual objects) of this distinction is carried out in the ventral stream and must be conveyed to the dorsal stream, an indirect indication of the interconnectivity between the streams. Alternatively for the sake of expediting action, neurons in the dorsal stream may 'learn' such distinctions (Grunewald, Linden et al. 1999) analogous to how, with a repeated task, V1 neurons become tuned to stimulus properties naturally only found in downstream areas (Zohary, Hillman et al. 1990).

The Posterior Parietal Cortex (PPC) as the endpoint of the dorsal stream is widely held to be a primary locus of sensorimotor transformations. For example, dorsal area 5, contains representations in both hand and eye-centered coordinate frames, indicating that information contributing to gaze position (such as head and trunk position) may be integrated there when movement plans are generated (Buneo 2002). Lesions of PPC generate symptoms that are related to the integration of vision and motor behavior. A PPC lesion results in optic ataxia (Milner and Goodale 1993), which is characterized by deficits in reaching with the contra-lesional hand, particularly to objects in contralesional space. Patients with right PPC lesions also are reported to have spatial neglect, which results in an attentional and intentional under representation of the left part of space (Karnath, Schenkel et al. 1991). Interestingly, even patients with severe neglect who are unaware of the neglected portion of space show strong indications of visual processing. For example in a case observed by Mashall and Halligan, a patient was repeateadly shown drawings of two houses identical on the non-neglected side with one house burning on the neglected half of space. The patient always claimed that the houses appeared identical to her, however she consistently chose the intact house as the one in which she would prefer to live, even though the placements of the two houses vertically was randomized (Marshall and Halligan 1988). [1]

[1] There has been much debate over how different cortical areas contribute to visual consciousness. Although parietal lesions produce deficits to visual awareness, patients with selective lesions to the

1.1.2 MIP and Parietal Reach Region

In order to distinguish between intention and attention in the IPS, Snyder, Batista, and Andersen conducted a study that held spatial attention constant but varied the effector (ocular motor vs. arm) (Snyder, Batista et al. 1997). They found anatomically distinct regions with differing neural activity depending on whether reaches or arm movements were executed to a cued region in space. The task, a memory reach, began with a color-coded cue (green for reaches and red for saccades) and concluded with a reach to one of 8 locations that were peripheral to a visual fixation point. An intervening period where no peripheral stimuli were present provided two analytical advantages. Firstly it separated visual responses from motor responses, and secondly it introduced a *planning* or memory period, where there was no movement, but the intention to reach was maintained. An analog of this area has also been identified in humans (Connolly 2003).

Further studies revealed that both areas maintained retinotopic coordinate frames (Batista, Buneo et al. 1999) as opposed to other physically and operationally proximal parietal areas that also included hand centered coordinate frames (Buneo 2002). Although other areas, such as the premotor cortex, are also modulated by the eye position (Boussaoud 1995), these results offer a unified coordinate frame that can be exploited in a prosthetic system. Although there has yet to be a comprehensive study comparing the variability of different cortical areas during reach plans, it can be inferred that independent of cortical region, prediction of movement can be improved by accounting for the variance given by eye position.

---

magnocellular pathways that project primarily to the ventral stream and MT/MST are known to have a condition termed *blindsight* due to their ability to produce effective visually guided movement in absence of any conscious vision. Goodale and Milner extensively use blindsighted patients to support theories of intentional vs. representational visual processing streams. Taken together, these lesion results suggest that visual consciousness may be the result of connectivity between the dorsal and ventral streams that is eliminated in both blindsight and neglect.

Prosthetic systems must maintain their efficacy in cases where movement is not possible, and/or atrophy has occurred. The evidence concerning motor cortical reorganization and atrophy depend on the nature of the patient's injury and what areas are studied. One fMRI study shows that dysmelic patients may have less ipsilateral reorganization (similar to normals) than amputees who have used stump-based prosthetics in (Cruz, Nunes et al. 2003), while another (Turner, Lee et al. 2001) found dramatic lower limb reorganization in the case of spinal cord injury and little reorganization in amputees. A study by a group conducting motor-cortex based BMIs found that the motor cortex in tetraplegics with spinal cord injury is not different from normals (Shoham, Halgren et al. 2001). With the exception of this final study it appears that motor cortex is at least partially susceptible to reorganization after spinal cord injury.

Due to its strong visual enervation and retinotopic organization, the parietal cortex is more likely to maintain a stable representation after injury. Assuming that neural implants cause minimal damage to healthy brain tissue, the best motor BMIs will likely include signals from multiple areas containing intentional information. It is likely that in the future BMIs will be customized to account for particular patients' conditions. For example, in the case of a stroke that disables the motor cortex, a PPC based BMI will be useful, whereas a patient with spinal cord injury may have a BMI that is localized according to fMRI data. For these reasons, a focused study of the efficacy of PRR as a basis for BMIs is justified.

1.2 Plasticity of the Dorsal Stream

The reach system is highly plastic, as has been demonstrated in adaptation experiments in which the visual feedback during reaching is perturbed with prisms (Held and Hein 1958). Errors detected for adaptation are in eye coordinates (registering the mismatch of the hand and reach target), and this would be the most natural coordinate frame in which to recalibrate reach plans. Knowing if PRR is plastic, and to what extent, is important for the design of a prosthetic for paralyzed patients, in order to optimize information capacity and performance of the system.

Although there have been relatively few studies that evaluate neural plasticity and adaptation on a cellular level, there is a wealth of behavioral data on human adaptation to prism induced mismatch of the visual and motor-proprioceptive aparatus. The type of adaptation induced in prism experiments is similar to the type of adaptation that may be necessary for patients with PRR based neural prostheses. Corrections made to place visual and motor/proprioceptive maps into register after prism adaptation may be similar to the calibrations that might be necessary to expand an initially limited sampling of motor representations by a neural prosthesis to one that is richer or more elaborate. Indeed, during prism adaptation in humans, the posterior parietal cortex is activated (Clower, Hoffman et al. 1996); theis has been confirmed by (Inoue, Kawashima et al. 1997). One of the potential benefits of PRR based neural prostheses is that, because of its endogenous plasticity and simple motor representations, plasticity in this region may be mediated without effortful control of neural activity. It has been shown that considerable effort and training time are required to induce plasticity in humans using a prosthesis placed in the motor cortex (Kennedy 1989).

Most electrophysiological and imaging studies of motor learning have focused on sequence learning. However, Wise et al. (Wise, Moody et al. 1998) performed experiments observing changes that occurred in MI, MII, and PMd during rapid adaptation to a number of different visuo-motor transforms. Most of the effects were changes in the gains of the response of the neurons. A preliminary study by Clower and Alexander found that during prism adaptation directional tuning of 66% of cells in area 5 of the parietal cortex changed their tuning curves in the direction of visually perceived limb position (Clower and Alexander 1997). This result suggests that cells in parietal areas such as area 5 and the closely located PRR are not exclusively locked to the kinesthetic properties of action, but rather adapt with the visual perception.

Psychophysical experiments have shown that both humans and non-human primates adapt rapidly to visual displacement induced by prism glasses in as few as 10 trials ((Helmholtz 1867); (Baizer, Kralj-Hans et al. 1999). The fact that repeated exposure of

humans to different prism displacements results in progressively faster adaptation (Welch, Bridgeman et al. 1993) may be a further advantage for maintaining the calibration of a neural prosthesis. The error signals in visuo-motor control that provide for on-line movement correction are largely visual in nature (Wallman and Fuchs 1998); (Ghez, Gordon et al. 1995), and they may even be impeded by proprioceptive information (Lajoie, Paillard et al. 1992). These data suggest that, provided sufficient visual information, the absence of proprioception will not hinder learning to use a prosthetic limb that is controlled without explicit motor effort. Finally, lesion studies have produced insights as to what the contributions of various areas are to differing types of adaptation. Evidence from lesion studies show that the cerebellum also plays a role in prism adaptation (Weiner, Hallett et al. 1983) (Gauthier, Hofferer et al. 1979). In non-human primates, inactivating ventral premotor areas eliminated adaptation (Kurata and Hoshi 1999). Most of these lesion studies do not control for active error-correction studies, as do Clower and colleagues. Thus, it is possible that the error-correction system and a cortical adaptation system function separately, with the former mediated by the posterior parietal cortex and the latter the cerebellum.

1.3 Other Successful Manuo-Motor BMIs

1.3.1 Open-Loop Experiments

Several open-loop experiments led up to the current closed-loop standard for multiple-single unit BMI development. For the purposes of this discussion, 'open-loop' experiments include all preliminary studies in which the subjects are simultaneously moving their limbs while an external device is controlled by the activity of single units. Although these studies did not in and of themselves generate BMIs that might be used by paralyzed patients, they do make an important advance beyond previous electrophysiology studies. This advance is in data treatment—rather than assessing only the significance of things like directional or spatial tuning in motor-related cells (often treated ergodically in post-hoc population analysis of single cell data), these studies rise to the challenge of producing control from cells in [soft] real-time. The first motor cortical open-loop experiment was implemented by Schwartz and colleagues (Schwartz 1996). Others have implemented a similar open-loop method using not only the motor

cortex but also the premotor (and parietal areas) (Hatsopoulos, Joshi et al. 2004) (Wessberg, Stambaugh et al. 2000). These types of studies [as well as offline movement reconstructions] have made steps toward assessing the information content of certain areas, as well as developing prediction algorithms for estimating trajectories. However, they have the drawback that they may inadvertently or intentionally incorporate information, such as somatosensory and proprioceptive signals, that is not likely to be available when a subject is unable to move the limb in question.

1.3.2 Closed-Loop BMIs

In the 1960s Fetz implemented a bio-feedback system in which monkeys were trained to change the activity of motor cortex neurons with visual feedback (Fetz 1969). This experiment demonstrated that cell firing could be influenced by visual feedback, and in repeated cases without concomitant muscular activity. There was even a case where two cells on the same electrode (which would typically fire similarly) were entrained to produce divergent activity.

Since then there have been several groups who have implemented closed-loop BMIs. The earliest of these experiments involved non-invasive methods using EEGs, research showing effective control dates to 1993 (Farwell and Donchin 1988). Recently these methods have advanced to levels that rival that of modern intercortical methods. These methods harness mu (8-12 Hz) and beta (18-25 Hz) signals recorded from scalp electrodes and train subjects to control the amplitude of these signals. The most recent advances in this research shows 45-70% accuracy in a 4 target task (McFarland, Sarnacki et al. 2003; Wolpaw and McFarland 2004) in well-trained subjects. This qualification, ‘well-trained’ ,is perhaps one of the most pressing issues when it comes to this type of research. Although many of the aforementioned EEG experiments were conduced in healthy subjects, there has been success with well-trained paralyzed patients, where binary control of EEG signals vary between 60 and 85% correct (Birbaumer, Kubler et al. 2000). Also, EEG based BMIs may be at the limit of performance, while BMIs based on cortically implanted electrodes has only recently begun to be developed. Controlling

EEG signals is an acquired ability, one that takes months of training to produce accurate results. This has also been found true in patients learning how to control unspecified single-unit signals not necessarily related to controlling factors such as arm movement (see below). Signals that are less abstract (such as those from the motor cortex and the parietal cortex) and related to spatial manipulation or reaching produce much faster results (even in subjects who are not aware of the task at hand).

The first intercortical experiment was conducted with a rat pressing a lever (Chapin, Moxon et al. 1999). The neural activity during a lever press was recorded and processed with a neural network whose output was the lever movement. Eventually the rats only had to initiate the movement, and the lever movement would be completed by the neural network. Though this was a modest advance it is significant in its attempt to close the feedback loop.

An electrode that used nerve growth factors and/or peripheral nerve tissue to stimulate neurite growth at the recording site was first implanted in amyotrophic lateral sclerosis (ALS) patients. Over a period of months the patients were able to produce rudimentary binary communication by learning what type of mental activity generated neural activity on the electrode, which was implanted in motor cortex (Kennedy 1989);(Kennedy, Bakay et al. 2000). Given that these patients were unable to communicate by any other means prior to the implantation and training, this small amount of communication was of great benefit.

Some direct demonstrations of BMI control have been based in the motor cortex (Taylor, Tillery et al. 2002);(Carmena J.M. 2003). With as few as 18 cells from M1, the full trajectory was predicted with an average of 49% tagets acquired out of 8 possible locations. The subjects, macaque monkeys, were placed in a ‘virtual reality’ style environment, where they had visual information about only the endpoint of their hand represented by a cursor in 3D space (training phase) or cursor continuously positioned at the predicted location of the hand (brain control phase) and the target. The authors attribute their success to the fact that the monkeys had continual feedback of the output of

the positioning algorithm and could thus adjust their neural activity accordingly, as well as a 'co-adaptive' algorithm, which adjusted to changes in tuning during the brain-control phase. Interestingly, in this experiment, the percent of targets hit correctly with the brain-controlled cursor was actually greater in the brain-controlled phase then when the cursor's position was predicted offline while the monkey was moving his arm. That is to say, visual feedback positively affected the accuracy (true even when the 'co-adaptive' features of the algorithm were eliminated). This was the first published experiment formally "closing the loop" using intracortical signals.[2] With training the EMGs of the monkeys' restrained arms diminished until there was no longer any sign of muscular activity, while the performance of the BMI continued to improve. Following this study, researchers developed a BMI that predicted position, velocity and grip force, using signals from M1, the dorsal premotor cortex (PMd), the supplementary motor area, the primary somatosensory cortex and the parietal cortex (Carmena J.M. 2003). These studies used up to 90 single units and multiunit activity from up to 175 sources. The performance was significantly above chance in all tasks (target acquisition, griping, and both) and in a task where only target acquisition was required (as above) the performance was over 90% correct. As in the above study, EMG recordings showed that arm movement was eliminated during the 'brain control' task. Results from training phases indicted that M1 cells generated the best predictions of all three movement components. The contribution of parietal cortex in these studies best correlated with gripping force, which may be reflective of the particular placement of the electrodes in the parietal cortex and/or the nature of the algorithm used in the study.

A BMI based in PMd and MIP was reported that did not decode trajectories (Musallam, Corneil et al. 2004), but rather the final targets of the reach. These researchers found that with as few as 8 MIP cells, a BMI could position a cursor to one of 6 randomly placed targets with 64.4% accuracy. With 16 PMd cells a cursor was correctly placed to one of 8 targets in 67.7% of trials. In the three experiments mentioned above, neurons changed

[2] Though others had been presented in informal or abridged formats (Meeker, D., S. Cao et al. (2001). Closed loop control of a neuroprosthetic. Society for Neuroscience, New Orleans, LA, Serruya MD, H. N., Paninski L, Fellows MR, Donoghue JP (2002). "Instant neural control of a movement signal." Nature **416**: 141-142.).

activity over the course of the experiments. The predicted value of the reward could also simultaneously be decoded from PRR activity including the type, magnitude, and probability.

1.4 Algorithms and Approaches

In limiting this discussion to algorithms and approaches that have been developed for BMIs we will omit several key developments in the decoding of sensory/motor input/output from neural activity. Instead this discussion will attempt to contextualize and motivate the approach taken in the experiments described in this dissertation.

There are two extremes of setting up the problem to be solved by movement prediction algorithms. At one extreme algorithms can be developed that map neural activity from a sampled space back to the same space. At the opposite end of the spectrum, the problem can be posed as solving the mapping from neural activity associated with a limited sample of movements to the entire range of movement space, including movement trajectory. In the first case the performance of the BMI is limited to how the space is sampled when the algorithm is developed and the robustness of that sampling. In the second, the performance of the BMI is limited by the efficacy of models that extrapolate from the sampled space to the complete space. This section discusses some examples of each of these models and describes how the first case might be extended to the complete space (and time) using non-neural data processing (and modeling of internal states based on neural data).

Perhaps one of the first offline models for mapping movement to a continuous representation of space is the population-vector (Georgopoulos 1988). Though this model is impractical in that it assumes radially symmetric cell tuning and uniform distributions of preferred direction, its vectoral approach of characterizing every sample of neural activity in time as a vector of the preferred direction with magnitude weighted by firing rate has been reproduced repeatedly in predictive algorithms. Slight modifications to this method correcting for the assumptions have produced successful

primate BMIs (Taylor, Tillery et al. 2002). Correcting for non-uniform direction distributions can be done optimally by regressing the population to the sampled movement directions to calculate each unit's preferred direction (Wessberg, Stambaugh et al. 2000) (rather than independently calculating the preferred direction). This optimization brings this version of movement prediction closer to the first case described above by optimizing to the sampled space.

A differing approach from the similar[3] linear filter and population vector methods is the Kalman filter. The advantage of using dynamic state models is that they model both the relationship of the neural activity to the movement and the dynamic process of the movement to itself. This approach has been applied in the reconstruction of rodent maze navigation using hippocampal place fields (Brown, Frank et al. 1998). More recently this has been applied in the context of neural prosthetics based in the motor cortex (Paninski, Fellows et al. 2004). This type of signal processing requires that all the features being modeled are sampled continuously, in opposition to the standard paradigm for arm movement, which has most frequently been a center-out reach. However, this study found that only the low frequency components of the movement were reproduced well with the Kalman filter, suggesting that despite the added complexity, low frequency components can be modeled by simpler means, and movement details may be derived not from neural signals but external systems.

The simplest method of the sampled-space to sampled-space decoder is to build models based on the neural activity and select the most likely source model from the neural activity, predicting the movement to the source location. One way of doing this is with a neural network that outputs a classification from the input of neural activity (Chapin, Moxon et al. 1999). Issacs and colleagues used a different classification approach, discretizing movement trajectories in the sampled-space and constructing a classifier that matches the neural activity to the most likely sub-trajectory (Isaacs, Weber et al. 2000). In this approach, the effort to extract trajectory information in this descretization process

[3] Similarities are described in Schwartz, A. B. (2004). "Cortical neural prosthetics." Annu Rev Neurosci **27**: 487-507.

creates many more source models and a dramatic increase in the number of samples that must be collected to build the models. Also, unless constraints are enforced, it is likely that discontinuous movement predictions due to signal variability will ensue. However, if only the endpoints of the movement are modeled, the number of source models decreases.

Most intracortical BMIs have focused on estimating trajectory using firing rate information from populations of single neurons, which complicates the estimation process dramatically and requires much more information than pure end-point estimation. In a review article Schwartz argues:

> Why is it so desirable to extract a trajectory signal from the brain? The trajectory of the end point contains natural characteristics of animate motion. Examples of these invariant features are the bell-shaped velocity profile of reaching movement and the two-thirds power law pertaining to drawing and handwriting. Although prosthetic devices can be effective without operating like natural limbs, the embodiment of these characteristics is desirable in terms of biomechanical compatibility with other body parts, ease of control, and aesthetics.

This argument, however does not take into account the success of robotic and artificial vision engineering. State-of-the-art robotics incorporates exactly the invariant features that Schwartz mentions. Although ideally the full trajectory would be available from the brain, current electrode technology limits the information that can be extracted from the brain, such that robotics and artificial vision are more likely to produce smooth movement than attempting to determine a continuous motion from neural signals alone.

For these reasons the experiments performed here apply a simple classification algorithm (maximum-likelihood) that maps neural activity to the sampled space directly. Models were only generated that map neural activity when a movement is being planned to the final target of that movement, minimizing the number of models that must be formed. This classification scheme is well suited to the nature of the task and the questions being

asked. The task isolates movement from planning and uses signals from an area that does not code the details of the movement trajectory, but the goal itself.

1.5 Adaptation and BMIs

The exact nature of the adaptation of neurons to BMI use is not clear, however it is clear that using a BMI does invoke change in neuron properties. In the study that included only 18 M1 neurons, the researchers found that preferred direction changed during the brain-control phase of the task. That is, not moving the arm changed the neural activity significantly, however the 'co-adaptive' algorithm used in this study compensated for these changes, which became systematic over days of BMI use, whose performance improved over the weeks that the experiments were conducted (Taylor, Tillery et al. 2002).

The interpretation of these results is challenged by Carmena et al. who found a reduction in tuning depth during the brain controlled phase of the task. They account for the differences in results by differences in ways that the tuning was measured, according to target locations in the first case and movement trajectories in the second. The authors cite differing views of M1 sensitivity to feedback from arm movement, but observe that even when the monkeys continue to move their arms during brain control there is a marginal reduction in cell tuning. "This suggests that directional tuning reflects neither movement dynamics nor abstract motor goals alone, but rather their combination." (Carmena, Lebedev et al. 2003). However, the two different tasks used in this study may also contribute to the differences observed. Musalam and colleagues also see improvement in performance over the course of BMI use, but they also observe that by manipulating expected reward value, performance was improved by up to 21%.

1.6 Local Field Potentials in Parietal Cortex

Informationally and physically, local field potentials (LFPs) lie between spikes and EEGs. EEGs, recorded non-invasively from scalp electrodes, integrate information over several centimeters of cortical space, through the barriers of the scalp, the skull, blood

vessels and dura. There may be LFPs generated by extracellular currents produced by local populations of cells; the sum of the EPSPs and IPSPs at the electrode tip (Buzsaki and Draguhn 2004). Previous studies in visual (Gray, Konig et al. 1989) and parietal cortex (Pesaran 2002) have shown that spikes are correlated with power in the gamma band of the LFPs. These studies have revealed that LFPs can be used to decode information such as the intended saccade direction with efficacy on par with that of spikes. Local field potentials possess many advantages that suggest that they should be included in the body of BMI research.

The local field potential may be collected from electrodes that are positioned in cortex whether or not spikes are present. Although the quality of both LFPs and spikes may degrade over time as scar tissue builds up around the electrode tip, one of the most active areas of research in electrode development is preventing this degradation in order to produce long-lasting signals (Vetter, Goodbody et al. 1999). How 'local' the potentials are depends on the impedance of the electrode. The greater the impedance, the smaller the radius around the electrode of currents that affect the signal; scar tissue effectively increases the impedance of the electrode, as most scar tissue is made up of astroglial cells, fatty tissue that creates a physical and electrical barrier between the electrode and local currents used to record LFPs as well as spikes (Moxon, Kalkhoran et al. 2004).

LFPs may also contain additional information about internal state over and above what may be present on spikes on the same electrode (Andersen, Musallam et al. 2004). This suggests that input information present in the EPSPs and IPSPs that may not be integrated by cells might still be represented in the field potentials.

2 Feasiblity Study: Finite State Machine Application for PRR Based BMIs

The prospect of helping disabled patients, by translating neural activity from the brain into control signals for prosthetic devices, has flourished in recent years. This is due largely to the successful demonstration of robotic arms guided by cortical activity measured during ongoing arm movements. To investigate how activity present before, or even without, natural arm movements might be used to control prosthetic devices we measured neural activity in parietal cortex while monkeys planned to reach to visually specified targets. Here we describe how such plan, or cognitive, activity from tens of parietal reach region neurons can specify when and where to move a prosthetic device. We propose that cognitive control signals may be well suited for use in a variety of prosthetic systems.

Neural activity previously recorded from the posterior parietal cortex was decoded using a finite state machine. In the task, monkeys planned to reach to visually specified targets. The state machine was able to predict where and when the animals intended to reach. Based on these results we propose that the intended movement signals from PRR may be well suited for neural prosthetic applications.

2.1 Introduction

The parietal reach region (PRR) of the posterior parietal cortex (PPC) is located at an early stage in the sensory-motor pathway. It is closely related to sensory areas, particularly visual areas, and projects to limb movement areas within the frontal lobe (Johnson, Ferraina et al. 1996); (Andersen, Snyder et al. 1997). Many properties of PRR make it an attractive source of plan activity to derive control signals for prosthetic systems (Shenoy, Kureshi et al. 1999); Meeker, (Cao et al. 2001). First, PRR plan activity is selective for arm movements, as opposed to eye movements, and persists until a reach is initiated (Snyder, Batista et al. 1997). The persistence of activity during planning does not require an actual movement; in essence this area codes the “thoughts” to move. This finding contrasts with motor cortex, where the activity is largely related to the execution of the movement (Maynard, Hatsopoulos et al. 1999). Second, PRR plan activity is

abstract, being represented in visual (eye-centered) coordinates, and the activity within the spatial representation shifts with each eye movement to remain spatially invariant (Batista, Buneo et al. 1999). Moreover, cells in this area also carry eye position information in the form of a modulation of the eye-centered response fields and thus the goals of movements can be read out in other coordinate frames as well (Cohen 2002). Finally, during sequential reaching to two memorized locations, PRR plan activity codes just the next intended reach (Batista and Andersen 2001). This simplifies the interpretation of activity in this region for prosthetic control since plan activity reflects the upcoming movement, not any or all planned movements. These properties suggest that intended movement activity from PRR may be well suited for generating high-level, cognitive control signals for prosthetic applications.

We report here the results of a computational investigation, using a database of PRR action-potential responses to explore how high-level, cognitive control signals can be estimated from plan activity using a finite state machine algorithm. This algorithm is well suited for the control external devices such as a robot limb or a computer (Wolpaw, Birbaumer et al. 2000); (Wessberg, Stambaugh et al. 2000); (Kennedy and Bakay 1998); (Kennedy, Bakay et al. 2000); (Serruya MD 2002); (Taylor, Tillery et al. 2002).

2.2 Methods

Single neuron recordings. The spike data recorded from PRR neurons was obtained from a previous study from our lab, and surgical and recording techniques for acquiring single-neuron action potentials have been described previously (Batista 1999);(Snyder, Batista et al. 1997). All protocols were approved by the Caltech Institutional Animal Care and Use Committee.

Data Analysis. We used maximum likelihood estimation, which is equivalent to Bayesian estimation with a uniform prior probability distribution, to estimate reach parameters. Our assumptions were Poisson spike statistics and statistical independence between cells, but explicit models of tuning to the various parameters were not assumed

(Zhang, Ginzburg et al. 1998). To reconstruct the planned reach direction, we defined the scalar x = (1, 2, …, 8) to be the reach direction and the vector **n** = (n_1, n_2, …, n_N) to be the spike count from each neuron (n_i) during a time interval (τ). Combining the expression for the conditional probability for the number of spikes **n** to occur given a plan to reach direction x with Bayes' rule yields the following expression for the conditional probability of x given **n**:

$$P(x \mid \mathbf{n}) = C(\tau, \mathbf{n})\, P(x) \left(\prod_{i=1}^{N} f_i(x)^{n_i} \right) \exp\left(-\tau \sum_{i=1}^{N} f_i(x) \right)$$

The normalization factor $C(\tau,\mathbf{n})$ ensures that the sum of the probabilities equals one. P(x) is the prior probability for reaches in each direction and is uniform by experimental design, and the mean firing rate of the i^{th} neuron while planning a reach to direction x is $f_i(x)$. The estimated reach direction, $\hat{x}$, was taken to be the one with the highest probability.

$$\hat{x} = \underset{x \in \{1,2,\ldots 8\}}{\operatorname{argmax}} \left(P(x \mid \mathbf{n}) \right)$$

Action potentials from 23 PRR neurons from monkey CKY and 41 PRR neurons from monkey DNT, were analyzed. All analyses yielded similar results for both animals. We used cross-validation techniques to assess the performance of this estimation process. For each repetition of the simulation, and in each of the eight possible reach directions, a random subset of the total number of cells was selected to avoid a cell sampling bias. One trial was selected randomly, from each of the selected cells, and set aside for use as test data. With the remaining trials from the selected cells, we calculated the average firing rates for each cell while planning to reach to each target. This mean was used as the rate parameter λ in Poisson distributions. The probability that a particular selection of test data belonged to each of the multidimensional distributions from each direction was assessed, and thus the most probable (i.e., decoded or predicted) reach direction was selected for each repetition in the given direction. This process was repeated 1000 times in each of the 8 reach directions and then normalized.

A similar procedure was used to estimate the response distributions for the time-course analyses, but with the following variations. After selection of the random subset of cells and the exclusion of a single random trial from each cell, the remaining trials were divided into 3 epochs: baseline, plan period, and pre-movement period (-600 to 0, 300 to 1000, and 1100 to 1350 ms, respectively, where 0 ms is the onset of the reach target, and reaches began directly after the pre-movement period ends). The trials from each direction, for each cell, and in each epoch were concatenated, and the data were sampled with 250 ms long moving windows with 50 ms time steps. The baseline epoch was concatenated across all directions. Additionally the plan epoch was also sampled using 500 ms windows rather than 250 ms windows. The mean of each epoch was used as the parameter for the single multidimensional Poisson distribution for the baseline period, and for each of the 8 multidimensional distributions for each direction in the 3 other epochs (the 250 ms sampled memory epoch, the 500 ms sampled memory epoch, and the pre-execution period).

Test-data firing rates were measured in 250ms windows, advanced 50 ms at each time step, through the duration of the test trial. The most probable condition (baseline, one of 8 plan directions, or one of 8 execution directions) was estimated independently in each time step as above.

2.3 Results

Data were analyzed from a previous study in which action potentials, eye movements, and push-button state were recorded while two monkeys performed a delayed center-out reaching task (Batista 1999). Figure 1A plots the response of a PRR neuron during repeated reaches to the memorized location of a flashed visual target. Three periods of neural activity are of particular interest: a *baseline* period preceding target presentation, a *plan* period following target presentation but preceding the reach cue, and a pre-movement or *go* period following the reach cue but preceding the onset of the arm movement. *Plan* and *go* period activity levels vary with the location of the flashed visual target, which specifies the goal of the arm movement. Figure 1B plots the average plan

period response of 41 neurons from PRR of the right hemisphere of one monkey (DNT), recorded sequentially while reaching in eight different directions. Most neurons are tuned for a particular goal direction, with other directions eliciting weaker plan period activity.

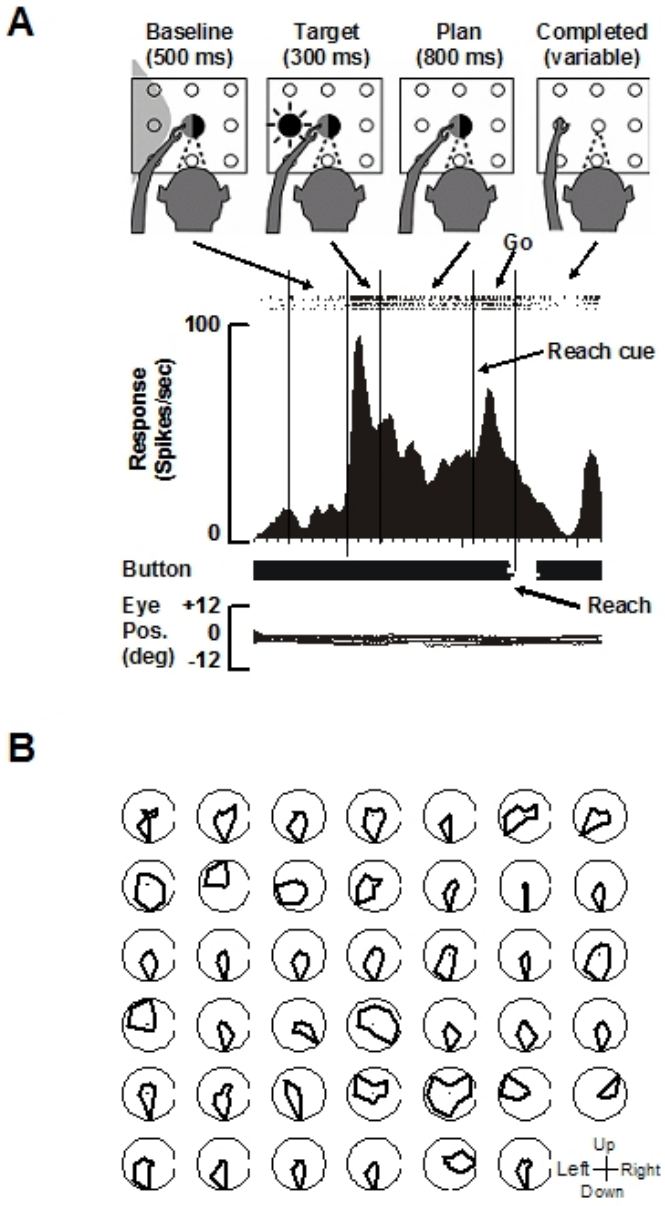

Figure 2-1 Parietal reach region (PRR) neural activity during the delayed, center-out reaching task. **(A)** The delayed, center-out reaching task consists of four stimulus/behavior periods (baseline, target presentation, plan and, go) as illustrated along the top of the panel. Icons depict eye (black semicircle and dashed lines) and hand positions (gray semicircles and arm), potential target locations (open circles), the neuron's region of maximum response or response field (large shaded region), and the location of the flashed target specifying the reach goal (black circle with emanating lines). The vertical line labeled reach cue indicates when the central eye and hand

LEDs are extinguished. Spike times are indicated as vertical lines in the trial-by-trial rasters (five rows at top of data) and the peri-stimulus time histogram (PSTH) represents the average response. Horizontal bars indicate when push buttons were depressed, with reach onset in a particular trial corresponding to when the bar vanishes. Eye position traces are shown at the bottom of the panel. Data are aligned to target onset (vertical line separating baseline and target presentation periods), which is when the reach goal first becomes known. This neuron from monkey DNT preferred downward reaches; the response field is illustrated to the left in the icon for illustrative convenience. **(B)** Directional tuning curves for each of the 41 neurons recorded in monkey DNT. Plan period activity (last 700 ms of the plan period) was averaged across all trials in each direction for each neuron to create the polar tuning curves. Plot angle corresponds to reach direction and plot magnitude indicates relative firing rate. Note the prevalence of downward directed tuning curves.

We wished to design a decode algorithm that can consistently and robustly determine from neural measurements alone: 1) when PRR is planning a reach, 2) when the animal intends to execute the planned movement, and, as already discussed, 3) in which direction the reach is being planned. Such capability is essential for future prostheses. Our approach to estimating these three parameters from PRR neural activity, and thereby generating high-level control signals, is illustrated in Figure 2. The top panel (A) illustrates the neural response from each neuron in the population throughout a representative (but simulated) delayed center-out reaching task. The vertical line labeled "t" represents the current time, which would also indicate the time of the most recent data if the prosthetic system were running in real time. Operating in real time, or causally, would mean that data to the right (ahead in time) of the "t" line would not be available. The middle panel (B) is termed the *classifier* and has two parts. The *direction classifier* uses neural data from the past 500 ms to estimate the probability that a reach is being planned to each of the eight directions, and the most probable reach direction is then selected. The *period classifier* uses neural data from the past 250 ms to estimate the probability that PRR is currently in a baseline, plan, or go period (see Figure 1), and the most probable class is then selected. The bottom panel in Figure 2 (C) is termed the *interpreter*. The *interpreter* must take in the series of baseline, plan, and go classifications, generated by the *period classifier* as time evolves, and determine when a reach should be executed. It must also take in where the reach should be directed from the *direction classifier* and finally issue the high-level control signal stating: reach here, reach now.

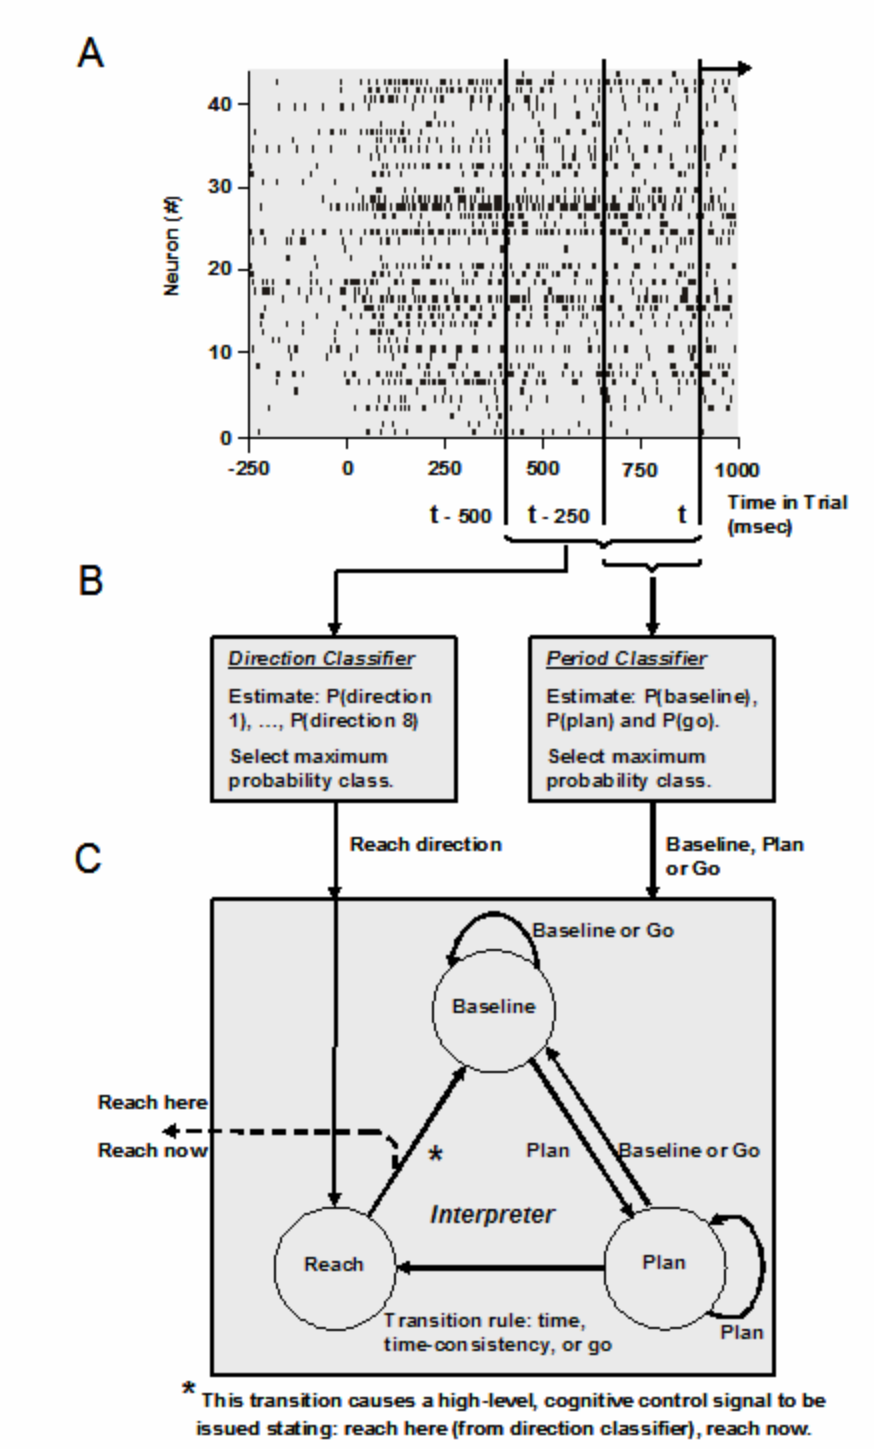

Figure 2-2 Computational architecture for generating high-level, cognitive control signals from PRR pre-movement, plan activity. See text for detailed description of function and operation. (A) Spike raster for each PRR neuron contributing to the control of the prosthetic device as a function of time in the delayed, center-out reach task. A single trial is illustrated and the visual target, specifying the eventual reach goal, occurs at 0 ms. The onset of arm movement occurs after 1100 ms (not shown). (B) Classifiers use

neural activity from finite-duration sliding analysis windows to estimate the direction of arm movement (direction classifier) and the current neural/behavioral period (period classifier). Both classifiers first calculate the probability of each class, and then select the most probable class for subsequent use. (C) The interpreter receives the stream of period classifications (i.e., baseline, plan, or go) from the period classifier and the stream of real direction classifications (e.g., downward reach) from the direction classifier. The interpreter consists of a finite state machine that transitions among three states (baseline, plan, and reach) according to the period classification at each time step. Three different rules for transitioning from the plan state to the reach state (time, time-consistency, and go) are considered. Once in the reach state, the interpreter always transitions back to the baseline state at the next time step in order to prepare for the next reach. During this transition a high-level, cognitive control signal is issued stating that a reach should occur immediately to the location specified by the direction classifier's current estimate. More sophisticated interpreters may include additional states and may use additional signals (e.g., band-limited LFP power) to govern transitions.

The *interpreter* starts in the baseline state and, as shown in Figure 2C, can transition to the plan state or return to the baseline state each time the *period classifier* issues another period classification. A baseline or go period classification keeps the *interpreter* in the baseline state, while a plan period classification advances the *interpreter* to the plan state. Once in the plan state, a baseline or go period classification will return the *interpreter* to the baseline state. The reason for this operating logic will become clear when we discuss below the possible rules for transitioning the *interpreter* from the plan state to the reach state. Once the reach state is achieved the *interpreter* automatically transitions back to the baseline state, and simultaneously issues a high-level, cognitive control signal commanding an immediate reach to the location given by the *goal classifier* (Figure 2C, asterisk).

The question of when to transition the *interpreter* from the plan state to the reach state, and subsequently triggering an arm movement, can be answered by considering the behavioral task instructions and go period classifications. We now summarize the logic of three different transition rules, as well as the results we obtained using these rules, to show how increasingly sophisticated rules can potentially improve prosthesis control performance.

Time transition rule. If the behavioral task instruction to the subject is simply, "plan a reach to a particular location for half a second," then a prosthetic system can safely execute an arm movement after detecting 500 ms of plan activity. In other words, the *interpreter* can transition from the plan state to the reach state when the *period classifier* issues 500 ms of contiguous plan classifications. Importantly, with this strategy the subject could abort an arm movement by ceasing to plan at any time before 500 ms or shift the reach target by simply changing his/her planned reach location before 500 ms has passed. We term this the "time" transition rule.

The *interpreter* begins in the baseline state and correctly stays in the baseline state during the phasic-response with cue onset. This time transition is introduced because the onset response is similar to the movement response and could otherwise result in an erroneous go period classification after only a brief (less than 500 msec) of planning. The *interpreter* then correctly enters the plan state and remains in this state, as long as the *period classifier* issues plan period classifications, until the minimum length of continuous plan classification (500 ms) is surpassed causing a transition to the reach state.

While this is the typical behavior with the time criterion transition rule, particularly with large neuron populations and for reaches to particular locations, this rule can err by failing to transition to the reach state before the end of the trial's experimental data or by executing a reach to the wrong goal location. Figure 3A shows the percent of trials achieving the reach state, and thus executing a reach, for a range of population sizes. Figure 3B indicates the percent of these trials that executed reaches in the correct direction for a range of population sizes. Ideally all trials would execute reaches, as all of our experimental data are from successful reach trials, and all trials would reach in the correct direction. Although this transition rule successfully executes reaches for most trials (Figure 3A), many of the reaches go in the wrong direction (Figure 3B). These errors are due to *direction classifier* misclassifications and are most likely caused by low signal to noise ratios. If errors were caused by drifts in plan or volition then the prediction

accuracy would not be expected to increase dramatically by adding more neurons to the estimate, as is seen in Figure 3B.

Time-consistency transition rule. A simple extension of the prior transition rule can address these concerns by adopting the conservative view that it is better not to execute a reach at all than to reach in the wrong direction. By adding the constraint that the *period classifier's* plan classifications must also specify a given goal direction throughout the required plan period (500 ms) we effectively impose a plan-stability requirement. We term this the "time-consistency" transition rule. Importantly, the *period classifier*, which employs a 250 ms sliding window, can also estimate goal location using response models and estimation methods analogous to those in the familiar *direction classifier*. Figure 3 also summarizes the performance of this transition rule. As expected, fewer trials now execute reaches (Figure 3A) but those that do tend to reach in the correct direction more often (Figure 3B).

Go transition rule. While the previous two transition rules perform well for certain applications, and importantly they do not rely on neural signals associated with movement execution, we would also like to be able to produce a larger absolute number of correct reaches. We can achieve this by replacing the previous stability constraint with a requirement that the *period classifier* issue a go period classification, after plan period classifications have been issued continuously for 500 ms, in order to transition from the plan state to the reach state. We term this the "go" transition rule. Using a neural "go signal" could afford the subject an additional opportunity to abort a planned reach by withholding the go command, or the possibility of reducing the length of the plan period on some trials. Figure 3 illustrates the performance. The period of time used by the *direction classifier* to estimate the reach direction, which is the 500 ms directly preceding the go period classification, tends to be slightly later than with the previous two transition rules. This is because the go period classification can occur up to several hundred milliseconds after the plan duration criterion has been met. This accounts for the increased percentage of reaches to the correct location (Figure 3B). This algorithm executes an intermediate number of reaches, as compared to the other two transition rules

(Figure 3A), with good performance arising from the readily detected and classified go activity.

Besides the increase in spike activity, another potential source of movement information is the local field potential (LFP). Previous investigations of frontal cortex have reported LFP oscillations (20-80 Hz band) during a pre-movement delay period, which cease around movement onset (Donoghue, Sanes et al. 1998), and Murthy and Fetz (Murthy and Fetz 1992) reported 25-35 Hz oscillations in sensorimotor cortex particularly during fine movements and focused attention. A recent study from our lab found similar results in the lateral intraparietal area in PPC in a 20 Hz centered band for saccadic eye movements (Pesaran 2002). To investigate the possibility that the PRR LFP may also be related to movement parameters, we recorded LFPs while a monkey looked and reached toward eight peripheral visual targets from a central starting position. In this additional animal a silicon micro-machined Utah Electrode Array with 22 active electrodes was implanted permanently within PRR for chronic recording. The LFP signals were filtered (15-25 Hz) to retain frequencies well modulated around the time of movement onset. The average power in this band is moderate around the time the central fixation and touch targets are illuminated, builds just before the peripheral targets specifying the saccade and reach goals become visible, and declines rapidly around the time of movement onset. Further examination revealed that power in this band is modulated by both saccadic eye movements and reaching arm movements. Reaching arm movements tend to modulate the power to a greater extent than do saccadic eye movements, with power being reduced to nearly zero directly after reach onsets.

A

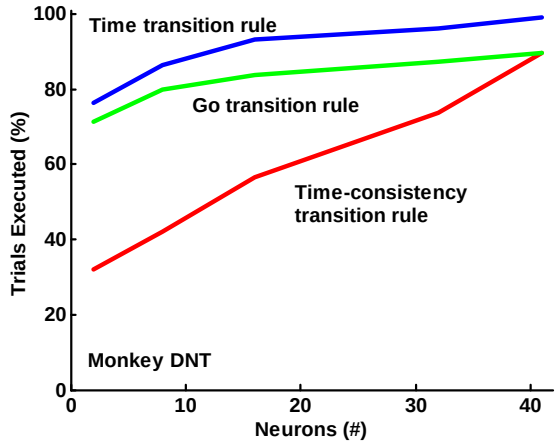
Time transition rule
Go transition rule
Time-consistency
transition rule
Monkey DNT
Trials Executed (%)
Neurons (#)
0
20
40
60
80
100
0
10
20
30
40

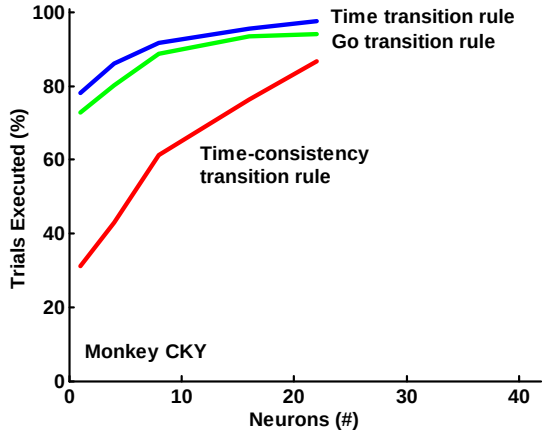
Time transition rule
Go transition rule
Time-consistency
transition rule
Monkey CKY
Trials Executed (%)
Neurons (#)
0
20
40
60
80
100
0
10
20
30
40

B

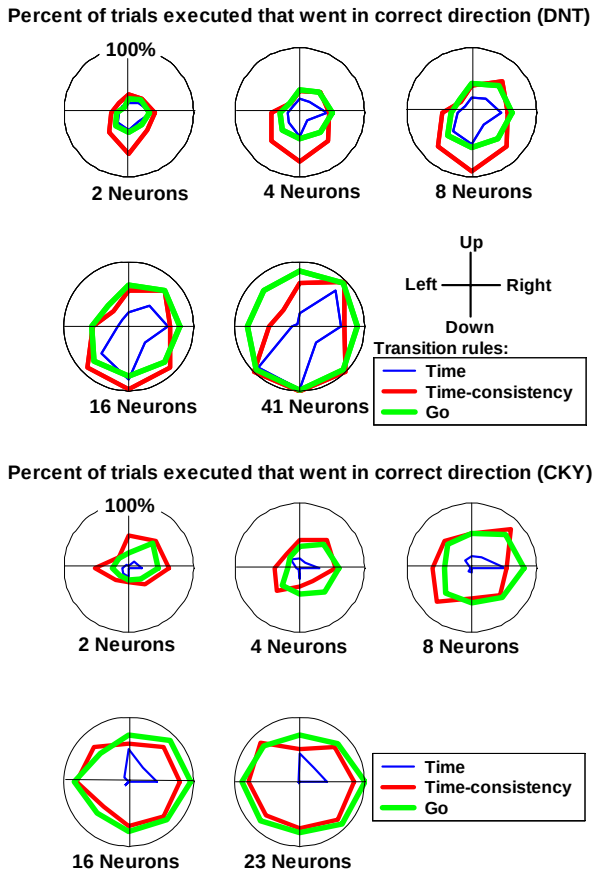

Figure 2-3 *Interpreter* performance characteristics. The *interpreter* was characterized separately, while using the time, time-consistency, and go transition rules (color coded). (A) Percent of trials achieving the *interpreter's* reach state, thereby triggering a reach, as a function of the number of neurons in the population. Perfect performance (100%) means that all trials executed a reach to some goal location, but not necessarily to the correct goal location. (B) Percent of trials that executed a reach to some goal location that *did* reach to the correct goal location. Perfect performance (100%), meaning that all trials executed went to the proper location, is plotted as a circle in all sub-panels. Each sub-panel shows performance for a different number of neurons in the analysis population. In both panels (A) and (B) neurons from animals DNT and CKY were used to generate the performance curves appearing to the left and right, respectively. Interpreter performance, including the relative performance of the three transition rules, was similar in both monkeys.

2.4 Discussion

Inspired by the considerable success of cochlear implants, tremor-control devices, and other neural-prosthetic systems aimed at *delivering* signals to the nervous system, research aimed at *reading out* neural signals for prosthetics applications has intensified in recent years , (Serruya MD 2002) (Taylor, Tillery et al. 2002); (Barinaga 1999) (Fetz 1999) (Mussa-Ivaldi 2000). While the concept of *translating* neural activity from the brain into control signals for prosthetic systems has existed for decades, substantial progress toward realizing such systems has been made only relatively recently. This progress has been fueled by advances in our understanding of neural coding, as well as by advances in forming stable electrical interfaces with neurons and computational technologies for processing neural signals in real time.

Despite these advances, the field of prosthetic systems that interface with the central nervous system is still in its infancy, and it is important to introduce new ideas about decoding movement parameters for possible use as prosthetic control signals since the full range of prosthetic applications is not yet known, and there does not yet exist a neural-prosthetic architecture that is optimal for all plausible prosthetic applications.

To explore the feasibility of using pre-movement neural signals from PRR to generate high-level cognitive control signals, we developed and tested the computational architecture presented in Figure 2. This part of an envisioned neural prosthetic system estimates, from PRR neural activity, when an arm movement is being planned (*period classifier*), the direction of the planned movement (*direction classifier*), and when the arm should move (*interpreter*). The resulting computations issue a cognitive control signal with two parts: reach here and reach now. Thus PRR contains sufficient signals to operate a neural prosthetic system.

To our knowledge, this is the first application of a state machine model to predict cognitive states using neural activity. Traditionally these techniques have been applied to

audiograms for the analysis of speech data. On the other hand, neural decoding algorithms have typically used an input-to-output structure that does not explicitly model internal dynamical states and the transitions between these states (Wessberg, Stambaugh et al. 2000) (Serruya MD 2002) (Taylor, Tillery et al. 2002). In many respects, neurophysiological experiments are designed with intuitive assumptions about states and state transitions in neural activity (for instance, the memory reach task used to generate the data in the current study.)

Possible attributes of PRR for prosthetics control

One open and central question is whether neural representations that are present during natural arm movements, and are employed in current prosthetic-arm research systems, remain completely intact following injury and disease or, alternatively, suffer at least some degeneration that would complicate their use in prosthetic control (Turner, Lee et al. 2001). Given that PRR is more closely linked to the visual system, and more distant from motor areas effected by paralysis, it is possible that it remains more intact following paralysis. Shoham and colleagues (Shoham, Halgren et al. 2001) have reported residual topography in motor cortex related to the will to move in partially paralysed patients. No doubt the most direct method of assessing the integrity of areas will come from cell recordings in paralyzed patients receiving prosthetic implants.

The parietal cortex is believed to participate naturally in ongoing visual-motor coordination and adaptation (Clower, Hoffman et al. 1996). Such cortical plasticity could help improve prosthetic system performance by continually, and quickly, adjusting for visual-prosthetic misalignments and by countering neural sampling biases, whether created by less than optimal surgical placement of electrodes or by a representational bias in cortex. A rapid and high degree of plasticity would enable patients to control a variety of devices including robotic devices that are very different from the human body, computers for communication, and autonomous vehicles. Plasticity will also be useful in allowing patients to effortlessly adjust to changes in the recordings that result from the usual small drifts of the electrodes in the brain. Kennedy and Bakay (Kennedy and

Bakay 1998) reported that several weeks were required to train a paralyzed patient to use motor cortex activity to move a cursor on a monitor. It will certainly be an advantage if the plasticity in PRR is more rapid.

In many forms of paralysis the patients also lose somatosensation. Somatosensory and proprioceptive feedback are important for error correction for motor behavior, as is vision. Since vision generally remains after injuries or diseases resulting in paralysis, and PRR is strongly and directly linked to visual cortex (Blatt, Andersen et al. 1990); (Johnson, Ferraina et al. 1996) , it is likely that PRR will still receive appropriate error signals for motor learning.

The cognitive quality of PRR activity has possible advantages. The persistence of planning activity, which does not require the execution of a movement, may be easily tapped in paralyzed patients who may still be able to activate this planning area, even though they cannot execute movements. At least in motor cortex, this planning-related activity, which precedes movement-related activity, appears not to be as robust in parietal cortex (Maynard, Hatsopoulos et al. 1999).

That the planning activity in PRR is in visual coordinates further emphasizes its cognitive nature (Batista 1999). Since both the retinal position of goals and current eye position are coded, in a separable form in PRR, the location of desired targets can be decoded in head centered coordinates. Recently we have examined the planning and execution signals for reaching to visually cued locations in area 5 of the somatosensory cortex (Buneo 2002). Interestingly, we find that the reaches are represented in a reference frame that is intermediate between eye and limb coordinates. The limb coordinate representation appears to be formed by subtracting the current location of the hand from the goal of a reach, with both locations represented in eye-coordinates. It is difficult to know what effects paralysis would have on such intermediate representations. Even if they remain intact, these intermediate representations are not separable into eye and limb centered representations and thus may be more difficult to decode than the representation in PRR. Experiments in the premotor areas suggest that a similar intermediate representation may

be present in premotor cortex (Cisek and Kalaska 1999); (Boussaoud and Bremmer 1999); (Mushiake, Tanatsugu et al. 1997). In motor cortex there is evidence for a limb centered coding of movement execution signals (Caminiti, Johnson et al. 1990); (Caminiti, Johnson et al. 1991). However, there has not yet been a systematic study of the motor cortex to determine if there is also a component of activity that is eye-centered; such a result would indicate that the motor cortex also codes in an intermediate reference frame.

Finally, the use of cognitive signals may reduce the number of neurons required for a given prosthetics application. This reduction is possible if relatively few, and high-level, parameters are estimated from the cognitive activity, and the signal to noise ratios are enhanced by averaging over the movement planning period.

Prosthetics system design using cognitive control signals

In order to produce complex movements we envision delivering high-level control signals to a reasonably sophisticated prosthetic controller capable of generating arm movement trajectories and capable of using inverse models of the prosthetic or electrically-stimulated arm to achieve the desired movement dynamics (Wolpert D.M. 2000). While the idea of such an intelligent prosthetic controller might sound fanciful at first, industrial robotics routinely combines state-of-the art machine vision and learning to achieve impressive levels of path planning, grip-force control, and safety. The patient would have the ability to plan an arm movement to an object, to have the controller guide the arm to that location, and perhaps to automatically grasp the object and, finally, the person could plan a subsequent arm movement to the desired location where the object could be released. Just as cognitive control signals could potentially cooperate with lower-level motor-cortical signals to further optimize control of arm-movement prostheses, cognitive control signals could also potentially contribute to existing communication-link systems (Kennedy, Bakay et al. 2000) due to the expected versatility of cognitive control signals as discussed below. Though using cognitive control signals may require more sophisticated prosthetic-system engineering than is currently

employed, these control signals may offer important advantages and reduce overall system complexity.

3 Rapid Plasticity in the Parietal Reach Region Using a Brain-Computer Interface.

Cells in the parietal reach region (PRR) of the posterior parietal cortex (PPC) of macaque monkeys encode the plans to make reach movements (Snyder 1997). A similar area has recently been identified in humans (Connolly, Goodale et al. 2000). We trained monkeys to use their intentions, decoded from PRR neurons in real-time with a probabilistic algorithm, to position a cursor on a computer monitor without actually making a reach movement. Approximately half of the cells recorded showed a rapid change in the directional differentiation of activity used to position the cursor. Control experiments show that the PRR activity does not predict the direction of intended eye movements or the direction of attention. The finding that the animals can quickly abstract their intended reach signals for cursor control, and rapidly learn to change this activity to improve performance in the task, suggests that PRR signals could be used to control neural prosthetic systems for paralyzed patients.

3.1 Results

Three rhesus monkeys were trained to perform a task that had three stages. First, reach directions that gave the highest and lowest activity for single PRR neurons were determined by having the animal perform memory-guided reaches in eight different directions. In this task a central fixation point appeared, and the monkey pressed that location on a touch screen and also fixated the light. A target was briefly flashed in one of the eight locations, and the animal memorized the location. A variable delay of 800 to 1200 msec followed, and it is during this period that the activity of PRR neurons reflect the intention of the animal to reach in a particular direction. After the delay the central light dimmed, and the animal reached to the remembered location of the flashed cue. In the second stage of the task the animal performed reaches in only two directions, the preferred and least preferred directions. Approximately 20 trials were collected in each

direction to build a model for predicting where the monkey plans to reach based on the firing rates during the delay period. In the third stage a cursor was enabled, which appeared at the termination of the delay period if the neural activity closely matched the model of one of the two targets, and the monkey did not move his arm. If the cursor did not appear, the central light dimmed, instructing the animal to reach to the remembered location of the initial cue. The progression of the behavioral task is shown in Figure 3-1.

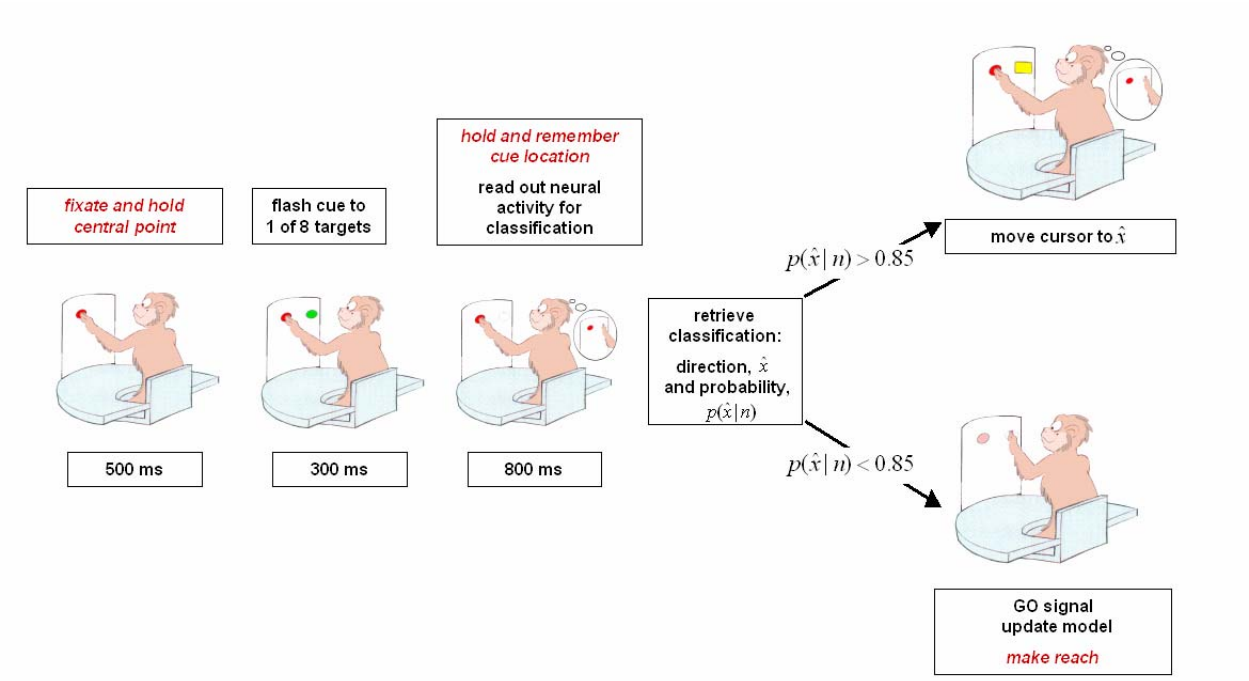

Figure 3-1 Task Progression. The sequence of trials begins with the monkey holding and fixating at the center of gaze. After a cue is flashed, the monkey waits for an 800-1200ms delay, neural activity during t=200ms-t=600ms are used for either building the database (downward fork) or moving the cursor (upward fork). All experiments begin exclusively with at least 20 trials of model building, after which the cursor is enabled and the probability that the neural activity belongs to either direction's model dictates the direction of the fork. If the activity closely matches the models of neural activity built during the database period (probability >0.85) then the cursor is moved; otherwise, the "go" signal is given, a reach is completed, and the model is updated.

The monkeys were able to correctly move the cursor based solely on their neural activity. Figure 3-2 shows how well the models of each neuron's activity predicted the monkeys' intentions. All trials, including those in which the neural activity was not deemed sufficiently close to either model and thus terminated with a reach rather than a cursor, were included in the analysis. Every trial was included in the analysis to avoid artificially inflating the results by only including trials that support the models, and to maintain chance performance at 50%. Average performance was 79.06%, and some cells demonstrated nearly perfect performance. Thus, given a small number of trials to generate models of neurons' selectivity, a binary movement plan can be interpreted from most cells with a much higher probability than chance, and without any arm movements.

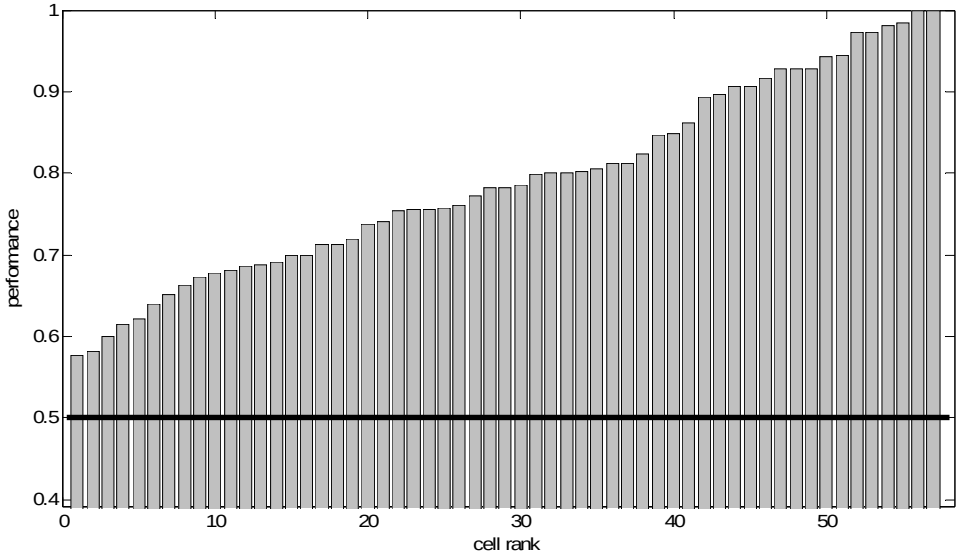

Figure 3-2 Rank ordering of performance across all cells collected. Fraction correct is on the vertical axis, and the rank of the cell is on the horizontal axis.

To assess the effect on the neural feedback on the firing rate we observe the cumulative difference in activity between the preferred and null directions. We define the sample sets $\Delta_{cursor}(t)$, the trial-by-trial difference between directions for the *t* trials after the first

cursor illumination, and Δ_{prior}, the trial-by-trial difference between directions for the first 20 trials prior to the first cursor appearance as:

$$\Delta_{cursor}(t) = \{f_p(t) - f_n(t), f_p(t-1) - f_n(t-1), \ldots f_p(0) - f_n(0)\}$$

and

$$\Delta_{prior} = \{f_p(-1) - f_n(-1), f_p(-2) - f_n(-2), \ldots f_p(-20) - f_n(-20)\}$$

where $f_p(t)$ is the firing rate in the preferred direction in the memory period on trial *t* after the first cursor introduction, and $f_n(t)$ is the firing rate in the null direction *t* trials after the first cursor introduction. A significant difference between $\Delta_{cursor}(t)$ and Δ_{prior} indicates a change in the relative firing rates from the 20 trials before the cursor appeared to the *t* trials after the cursor first appeared.

In 33% of the cells (19/57) enabling the cursor produced a change in the cells' responses that significantly increased the difference between the activity in the preferred and null directions ($p<0.05$, kruskal-wallis). Two cells significantly decreased this difference, while the remainder maintained the same relative firing rates before and after cursor appearance. None of the 57 cells showed a significant regression (ANOVA, $p<0.05$) in the relative firing rate during the time period prior to the first cursor's appearance, including some cells that had up to 60 trials preceding the first cursor appearance.

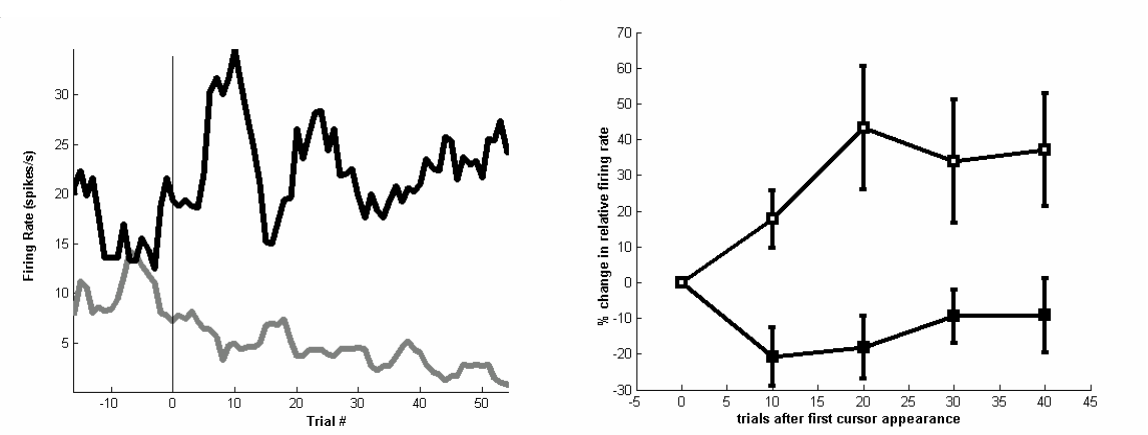

Figure 3-3 (A) A single cell's firing rate over the course of the feedback experiment. The cursor was first moved at trial 0, indicated by the vertical line. The gray line indicates the firing rate in the preferred direction and the black line shows the firing rate in the null direction. (B) The cumulative percent change in the firing rate compared to the model building period for all cells from CKY and GDY showing a significant increase in relative firing rates before and after cursor movement (11/23, mean ± S.E.).

Figure 3-3A shows an example of this learning in a cell that had a significant increase in selectivity only 20 trials after its activity was first used to position the cursor. In this case the change was due to both an increase in firing in the preferred direction and a decrease in firing in the null direction. Figure 3-3B shows the cumulative change in activity after 10, 20, 30, and 40 trials for all cells with significant relative changes in firing rates. The change is a result of both an increase in the preferred direction and a decrease in the null direction.

Not all cells that demonstrated this increased differentiation showed the effect in the first 20 trials; in some cases the effect was not significant until 50 trials in each direction had been executed. With one exception, all the cells showing the effect maintained the difference in activity for the duration of the experiment. This result indicates that behaviorally relevant feedback, such as cursor control, can result in rapid modification of neural activity in PRR, and that not all cells show the effect with the same time-course.

We also conducted a further experiment, and offline analysis, to determine whether attention is the determining factor in the control of the cursor. In monkey CKY a memory task was performed in which green flashed targets instructed him to plan and execute reaches and red flashed targets instructed him to plan and execute saccades. Using a population version of the prediction method used in the cursor control task, we considered how well a single trial of delay period activity from each of 41 cells could be used to predict both the direction and type of the intended movement. If the activity being used for cursor control was related solely to attention, the type of movement, saccade or reach, would not affect the performance of predicting the direction of the planned movement.

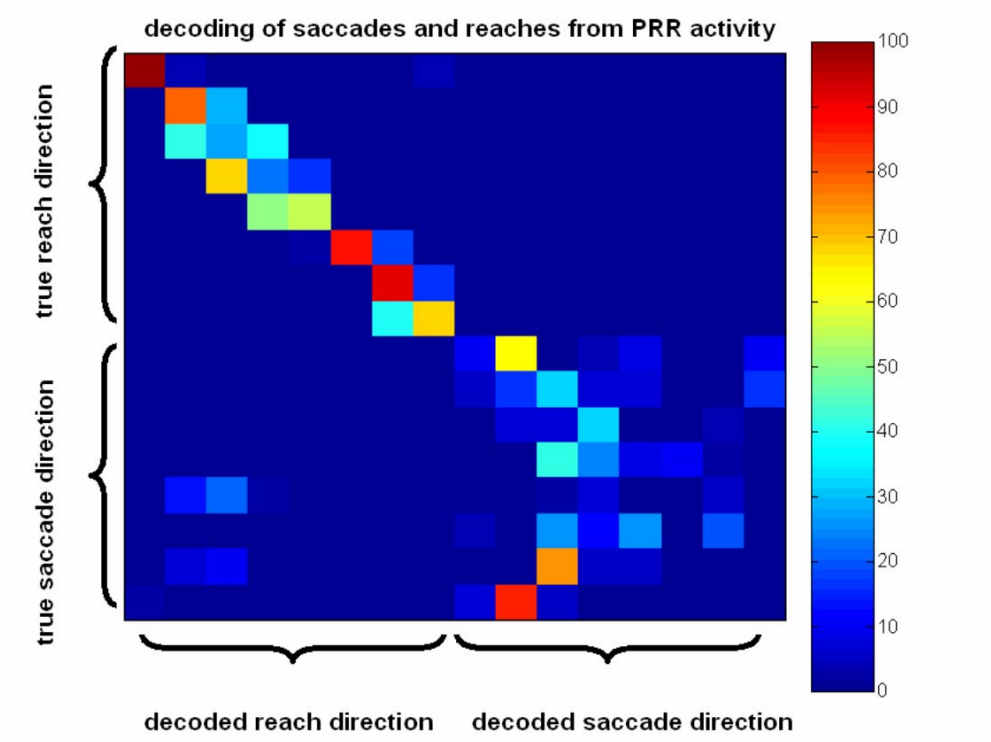

Figure 3-4 Confusion Matrix for decoded reaches and saccades.

The results of the population analysis are displayed in the 16 X 16 confusion matrix in Figure 3. The 8 X 8 upper left quadrant shows delay activity prior to reaches classified as

reaches. The fact that most of the cells fall into the small squares along the diagonal indicates that the planning activity is classified for the correct directions in reach trials. The lower right quadrant shows that the activity during the delay period in PRR leading up to saccades is not a good predictor of intended saccade direction. Since the animal likely attended to the locations for reaches and saccades, and there is an obligatory shift in attention toward the goal of a saccade (Kowler, van der Steen et al. 1984), it is thus unlikely that attentional effects on PRR activity can be used as a predictor of direction. The upper right quadrant displays reaches misclassified as saccades, which never occurs, and the lower left quadrant saccades misclassified as reaches, which rarely occur.

In 7 of the 57 cells tested for cursor control (all from monkey CKY) a block of delayed saccade trials was performed at the end of the experiment using the same two locations as the prior cursor control block. Analysis of the cursor control and saccade data, using 4 X 4 confusion matrices for the two directions and 2 movement types produced similar results to the 8 direction experiments with reach directions being much better classified than saccade directions.

Sensory neurons have been demonstrated to alter responsiveness with repeated exposure to a stimulus (Zohary, Celebrini et al. 1994) (Recanzone, Merzenich et al. 1992) (Schoups, Vogels et al. 2001) (Crist, Li et al. 2001). We did not find that repetition of trials prior to enabling the cursor increased the difference between the firing rates in the two directions. However, in several cases enabling the cursor resulted in immediate changes in cells' firing rates for preferred and null directions. This difference may be due to the differing nature of the task, which is not a threshold detection task, but an action-planning task. Additionally, the feedback that the monkey is given is directly related to the neural activity, rather than the perception of the stimulus.

Other studies are suggestive that PPC plays a role in rapid sensory-motor adaptations and associations that may be similar to the type of learning we observe in the cursor control task. A functional imaging study (Clower, Hoffman et al. 1996) indicates that changes occur in the posterior parietal cortex when reaching is adapted to compensate for distortions in visual feedback introduced with prisms. The learning of arbitrary associations between movement directions and visual stimuli produces rapid changes in neural activity in pre-motor cortex that lags the even more rapid behavioral learning (Mitz and Wise 1987). This result suggests that the locus of learning is antecedent to premotor cortex, and perhaps includes PPC. Cells in area LIP of monkeys trained to saccade to the locations of sounds have auditory receptive fields (Grunewald, Linden et al. 1999). LIP is a visual-oculomotor area adjacent to PRR that lacks auditory response fields in untrained animals (Grunewald, Linden et al. 1999), implying that it may be involved in learning associations between sound locations and saccades, similar to learning for arbitrary associations.

3.2 Discussion

This was the first study, to our knowledge, to directly demonstrate that a thought to move alone can be used to build models for controlling external indicators in a spatial task. To date other experiments in this vein have used peri-movement activity from motor cortex to build models later used to control external indicators (Fetz 1969; Chapin, Moxon et al. 1999; Serruya MD 2002; Taylor, Tillery et al. 2002). This distinction may be important if prosthetic applications are considered, due to the fact that movement is not possible in paralyzed patients, and could potentially lead to difficulty in model construction. However, in one of these experiments, restraining the animals limb did ultimately led to effective control without EMG indication of isometric muscle activity (Fetz 1969). We used neural activity in a cognitive, planning period in which the limb's position is

unrestrained, but maintained at one position without movement. We also control for gaze variables, and visual inputs, which have been shown to modulate activity in both the motor cortex (Baker, Donoghue et al. 1999) (Wannier, Maier et al. 1989) and the parietal cortex (Batista 1999). These additional measures of task control serve to show that the signals being used to position the cursor are truly related to the movement plan. Interestingly, the memory reach task used here evokes strong planning activity between the offset of the cue and the reach or cursor positioning. Similar memory reach tasks used in motor cortex experiments have failed to strongly activate most motor cortex neurons (Maynard, Hatsopoulos et al. 1999) (Smyrnis, Taira et al. 1992) (Crammond and Kalaska 2000). Thus the specialization of PRR for movement planning may be an important attribute for prosthetic applications..

The sensory-neural task using a brain-computer interface represents an interesting new paradigm for studying learning and other cognitive functions. The task acquisition and adaptation we observe in these experiments occurs on the same time scale as phenomenon such as prism adaptation, as opposed to the weeks required to modify neural activity observed in similar binary tasks in human EEG and motor cortex (Wolpaw, Birbaumer et al. 2000) (Kennedy and Bakay 1998). In these latter studies the subjects may have difficulty identifying what mental activity is necessary to induce external control. These observations suggest that the sensory-neural task may provide clues about the natural functions of a cortical area based on the ease with which the animal can control particular variables. A second feature of this technique is that the degree of plasticity capable in different cortical areas can be easily assessed and compared. For instance, it can be used to determine if PRR neurons only learn continuous distortions between sensory and motor maps, or can they learn arbitrary mappings, and how does PRR plasticity compare to motor cortex using the same tasks. This technique can also directly test which feedback signals are used by different cortical areas for learning. For

instance in the current study we have dissociated visual feedback from proprioceptive feedback. Thus, this technique offers a new electrophysiological method, that can be included with the currently used paradigms of examining correlations between behaviors and neural activity, for examining the functions to a cortical area.

This alternate "sensory-neural" paradigm for investigating learning suggests that neural activity in PRR may visually encode a high-level representation of a target for action, based on visual inputs. Other experiments in real-time decoding from cell populations in the motor cortex (closed loop) (Taylor, Tillery et al. 2002); (Carmena J.M. 2003); (Serruya MD 2002) , and the motor, premotor and parietal cortices (open loop) (Wessberg, Stambaugh et al. 2000) have suggested that using this type of real-time decoding is feasible for controlling prosthetic devices. Particularly when scaled up to the level of the populations of cells (Shenoy 2003), the high-level, plastic representation present in PRR could be useful for a brain-computer interface to operate a neural prosthetic for paralyzed patients.

3.3 Methods

Electrophysiology. Surgical and recording techniques have been described previously (Snyder 1997) (Batista 1999). The eye positions of one eye in animals CKY and DNT were monitored with the scleral search coil technique and in animal GDY with an optical eye tracker (ISCAN). The experiments were performed in a dark chamber, and the monkeys were observed using an infrared camera and infrared light source.

Behavioral Tasks. Two trial types were used, memory reach trials and cursor control trials. The monkeys were trained to reach to a touch-sensitive screen positioned at the center of gaze. Both trial types begin identically, with central green and red targets (squares 1° in width and height) being illuminated in the center of a touch-sensitive

screen. The monkey fixates the red light and places his hand over the green light. After 500 ms a green cue is flashed for 300 ms in one of 8 locations 18° from, and evenly distributed around, the center. The monkey maintains his eye and arm position for a randomized delay period (800-1200 ms), after which either the central targets dim signaling that the reach may be executed (delayed reach trials), or the cursor is moved to the decoded reach location (cursor control trials). If the reach or predicted reach was within 8° of the correct location within 1000 ms, a juice reward was delivered. For monkey CKY, the monkeys' eyes were required to maintain gaze within 8° of the fixation spot during the memory period but were free to move during the reach. Monkeys DNT and GDY were required to maintain their eye position within 7° of the central target during the entire task. The delay period activity is similar in PRR whether or not the eyes are moved at the end of the trial[i]. For the attention control task, the monkeys performed only a reach or a saccade movement when the central targets were extinguished, depending on whether the initial cue was red (saccade) or green (reach). In seven cells, 10-20 saccade trials in either direction were executed following the cursor control task.

Task Progression. For each isolated cell, the delay reach task was performed up to 20 times in each of 8 directions. Cells without tuning for reaches were not further recorded (20% in PRR). The two most separable directions were chosen by selection of the two directions with the maximum Fisher Criterion. Upon choosing these directions, the task was repeated in only these two directions, and a database of activity was constructed for use in the online decoding. After the database was constructed (approximately 20-30 trials in each direction), the decoder was enabled. The decoder returned both an estimated direction and a probability that the reach was planned in that direction ($\hat{x}$ and $P(\hat{x}|n)$, respectively). In order to keep the monkey engaged in the task, the cursor was only moved if the $P(\hat{x}|n)$ exceeded 0.85 (67.2% of all trials on average), otherwise a delay reach trial was carried out, and the database was updated.

Prediction. Maximum likelihood estimation, which is equivalent to Bayesian estimation with a uniform prior probability distribution, was used to estimate reach parameters. We assumed Poisson spike statistics but no explicit models of tuning parameters were assumed. We defined x = {1, 2} to be the possible reach directions and n to be the spike count from the neuron during a 600 ms time interval beginning 200 ms after the cue was extinguished. Combining the expression for the conditional probability for the number of spikes n to occur given a plan to reach in direction x with Bayes' rule and a Poisson assumption for the mean firing rate distribution yields the following expression for the conditional probability of x given n:

$$P(x \mid n) = C(n)\, P(x)\, f(x) \exp\big(-f(x)\big)$$

C(n) is a normalization factor and P(x) is the prior probability for reaches in each direction, which is uniform in this study. The mean firing rate of the neuron while planning a reach to direction x is f(x). The estimated reach direction, $\hat{x}$, is the one with the highest probability.

$$\hat{x} = \underset{x \in \{1,2\}}{\operatorname{argmax}}\big(P(x \mid n)\big)$$

The $P(\hat{x} \mid n)$ and $\hat{x}$ were then conveyed to the behavioral display as described above.

A population version of this method was used for the comparison between saccades and reaches, 16 (8 reach directions and 8 saccade directions), rather than 2 possible models,

were created, and *n* is a vector of neural activity across a population of cells rather than a single cell's activity. Cross-validated Monte-Carlo simulations were used to generate performance percentages. For further details, see Shenoy et al. 2002. This saccade/reach comparison was also performed offline on 7 of the cells used in this study, in this case x={reach 1, reach 2, saccade 1, saccade 2}, and *n* was the firing of the single cell in question during the delay period prior to the reach or the saccade in either direction.

Statistics. In order to assess changes in the firing rate over trials, the firing rate in each direction as a function of trial number aligned to the first cursor appearance was calculated. The resulting functions were subtracted from one another, and this function was tested for significant slope in a linear regression in the pre-cursor time period (ANOVA, $p<0.05$). For the remainder of the statistics we used Kruskal-Wallis test (Mann-Whitney) due to the non-gaussian distribution of our samples and the unequal sample sizes being tested.

To assess if there was a change in the cells' relative firing rates following cursor onset, we tested for a difference between $\Delta_{cursor}(t)$ and Δ_{prior} (Kruskal-Wallis, $p<0.05$). We compared the first *t* trials after cursor onset to the first 20 trials prior to the cursor's appearance (n={10,20,30. . .50}) in order to determine at which point the cumulative difference became significant. We did not choose to use a linear regression to assess this change because many of the cells adapted within only a few trials.

4 Using Local Field Potentials for Brain Machine Interfaces

The single largest challenge in the development of a reliable brain-machine interface (BMI) is in accessing enough information from the brain to accurately control an external device. This study explores the use of local field potentials (LFPs) as an alternative to the standard sorted-spike method of neural control. In 33 serially executed experiments we intentionally positioned an electrode at a site in the medial interparietal cortex (MIP) where no isolatable spikes were present while a monkey performed reaches to remembered locations. We gathered a database of LFP signals. Each database sample was comprised of a feature vector consisting of the power present in ten 10 Hz bands from 1-100Hz during reach planning. After the database period, the BMI was enabled, and a cursor was positioned to the location predicted by a maximum-likelihood estimate of the desired reach while the monkey maintained his original fixed hand position. The following results were obtained: (1) The LFP-driven BMI performed with 69.0% accuracy in a binary task, 16% worse than a spike-driven BMI performed in the same task. This is not as high as would have been predicted from former offline experiments using LFPs from locations where spikes were optimally isolated. (2) Offline analysis showed that the BMI did show a trend towards improvement in performance within each session, which may be a reflection of improved neural discrimination with practice or increased motivation for the less difficult cursor-control task. (3) The prediction using LFPs was weaker than previous offline studies that optimally isolated spikes. A postliminary analysis revealed some of the causes for the weaker than expected performance. A dataset comprised of simultaneous recordings from 48 electrodes implanted in PRR bearing various levels of spike amplitudes and LFP tuning showed a correlation between spike amplitude (maximized in the former offline experiments) and LFP tuning. (4) Canonical Variate Analysis (CVA) indicated that the plan information was most dominantly present in the 20-40 Hz bands, while visual cues for reaching were predominantly present in 50 Hz bands, and visual cues for saccades were predominantly encoded in 90-100Hz bands. Thus the performance of an LFP-driven BMI can be improved by selecting the optimal frequency bands, and eliminating frequencies that carry information that is not related to direction discrimination from the classification.

Although the LFP-driven BMI did not perform as well as the spike-driven BMI, the LFPs are still effective for neural control, while having several practical advantages that can improve the performance and extend the life of electrode implants.

4.1 Introduction

Brain Machine Interfaces have been largely based on EEGs or single-unit activity. The EEGs, though non-invasive and relatively easy to collect, have lower information content than do single-units (Vaughan, Wolpaw et al. 1996). However, BMIs that are based on spikes require surgically implanted electrodes that do not always have a high yield of single units (Nordhausen, Maynard et al. 1996; Williams, Rennaker et al. 1999; Nicolelis, Dimitrov et al. 2003) and when present may have a limited lifespan due to astroglial scarring at electrode tips (Rousche and Normann 1998). Although other [unpublished] offline studies report effective performance of LFPs for prediction of reach direction, these were conducted while electrodes were proximal to directionally tuned spiking neurons. In order to justify the risks associated with neurosurgery, it may be helpful to also take advantage of the LFP signals that are present on electrodes even when they are not close to neurons and only LFP signals are present in the recordings.

Informationally, the LFPs lie between EEGs and single-units, the signals are thought to indicate the combination of the EPSPs and IPSPs near the electrode tip, Buzsaki proposes that LFP activity is largely the reflection of inhibition that entrains post-synaptic potentials (Buzsaki and Draguhn 2004). This study demonstrates the use of LFPs for a BMI where a monkey positions a cursor to a cued location with a focus on investigation of the properties of the LFPs that contribute to the performance of the BMI.

4.2 Methods

Local Field Potential Signals. The monkey was prepared with microdrive chambers positioned over medial intraparietal area (MIP). At the start of each experiment a platinum-iridium electrode was positioned to a location with cells that were responsive to the memory reach task. The electrode impedance was 1-1.2MΩ. Following this initial positioning, the electrode was advanced a minimum of 200um to a location where no isolatable spiking signals were present. The local field potentials were recorded with a lowpass filter [1-120Hz]. Spectral estimates of the data used for the BMI were calculated using Thompson's multitaper estimate (Thompson 1982). Ten features were extracted from each estimate, corresponding to frequency bands centered at 5,15…85,95 Hz. These feature vectors were the inputs to the BMI.

Reach Task. The monkey was trained on the memory reach task. The targets were separated by 45° on a touch screen illuminated from behind with a projector. The center of the circle was located 30 cm from the monkey's eyes, the 8 targets were located 18° degrees of visual angle from the central location.

The task progression begins with the monkey reaching to and visually fixating a light illuminated at the center of the workspace. The monkey maintains fixation while a visual cue to one of eight possible locations is illuminated for 300 ms. This fixation position is maintained for a planning period of 800-1200ms and ends with the execution of a reach to the desired location only after the central light is extinguished. Visual fixation must be maintained at the center for the duration of the reach. Correct completion of the task results in a juice reward.

This task was repeated a minimum of 20 times in each of these possible directions. LFP data was collected during the time period 300-800 ms after the end of the cue period. The two directions with the greatest discriminability were selected for the binary BMI task.

BMI Task. After repeating the reach task a minimum of 10 additional times to compose a training dataset from the memory period in each of the two selected directions, the BMI was engaged. Cursor trials positioned a cursor to the predicted location of the reach at the end of the memory period. The prediction had to have an 85% likelihood score in order to result in a cursor placement, otherwise the go signal was given, and the monkey performed a reach. A correct placement was followed by a juice reward, an incorrect replacement terminated the trials. These trials were interleaved with reaching trials that did not depend on the likelihood score, in order to maintain the monkey's interest in the task and provide data that was not dependent on the BMI.

Saccade Task. For several sites, we also interleaved saccades to the selected target directions with the reaches. The saccadic task was identical to the reach task, differentiated by the color of the cue, and a saccade, rather than a reach at step 4 above. The monkey must maintain constant hand position during the saccade.

Direction Selection and distance metric. We used the mahalanobis distance as a selection metric to pick the best directions and as a measure of discriminability. The distance between the vector x_r and the source distribution for x_s is given by

$$d_{xs} = (x_r - x_s)\, C^{-1}\, (x_r - x_s)'$$

where C is the covariance matrix. This gives a way of estimating distance that takes into account the variance of the distributions, minimizing the possibility that we select two directions for discrimination that are spatially distant but having such high variances that they are not useful for discrimination.

Prediction Algorithm. We used the maximum-likelihood estimate of the data. The frequency bands were considered independent, giving a diagonal correlation matrix, C. For M repetitions and 10 features (n) the Gaussian probability distribution function is:

$$g_{[m,C]} = \frac{1}{(\sqrt{2})^n \sqrt{\det(C)}} e^{\frac{-(x-m)^T C^{-1}(x-m)}{2}}$$

Applying Bayes theorem, a probability for each of the two possible directions ($C_{k=1,2}$) is given by

$$P(C_k \mid x) = \frac{P(x \mid C_k) P(C_k)}{P(x)}$$

If the maximum P given the data vector (***x***) was >0.85, the cursor was positioned to C_k.

Canonical Anlaysis. See (Campbell and Atchley, 1981) for further details.

Given *g* classes, each with n_g training objects such that $x_{ki} = (x_{i,ki},...x_{M,ki})$ for all $k=1,..,g$ and $i=1,...,n_k$ then a canonical variate analysis forms a linear combination, $y_{ki} = c^T x_{ki}$, of the input attributes such that the ratio of the between-groups sum of squares,

$$SS_B = \sum_{k=1}^{g} n_k (\bar{y}_k - \bar{y}_T)^2$$

$$SS_B = \sum_{k=1}^{g} n_k (\bar{y}_k - \bar{y}_T)^2$$

,

and the within-groups sum of squares,

$$SS_W = \sum_{k=1}^{g} \sum_{i=1}^{n_k} (y_{ki} - \bar{y}_k)^2$$

,

where $\bar{y}_k = \frac{1}{n} \sum_{i=1}^{n_k} y_{ki}$ is the mean of the *k*-th class, $\bar{y}_T = \frac{1}{n_T} \sum_{k=1}^{g} n_k \bar{y}_k$ is the overall mean and $n_T = \sum_{k=1}^{g} n_k$ is the total number of training objects.

Substituting $y_{ki} = c^T x_{ki}$ gives

$$SS_W = c^T \sum_{k=1}^{g} \sum_{i=1}^{n_k} (x_{ki} - \bar{x}_T)(x_{ki} - \bar{x}_T)^T c$$
$$= c^T W c$$

Given *p* frequency bands, there are *h*=min(*p, g-1*) canonical vectors with non-zero canonical roots. If $C=(c_1,...,_h)$ and $F=diag(f_1,...,f_h)$ then the eigenanalysis becomes *BC=WCF*.

(From Evans, F. An investigation into the use of maximum likelihood classifiers, decision trees, neural networks, and conditional probabilistic networks for mapping and predicting salinity. (1998).)

We used fixed width windows to calculate the canonical variables and eigenvectors for the cue and memory periods in the static analyses of optimal memory and cue period discrimination, and sliding windows overlapping by 200ms with a width of 300 ms for the continuous analysis.

We chose to use this method to produce an optimal estimation of the reach direction and determine which frequencies contribute to this estimation.

4.3 Results

4.3.1 Performance of LFP-Driven BMI

We evaluated the performance of an LFP driven BMI. When considering exclusively the trials where the monkey positioned a cursor based on neural activity, the mean performance was 69.0±13.24% correct, with chance level of 50%. Figure 4-1 shows a histogram of the performance for each site. On average, there were 62 cursor positioning trials for each site, with reaching trials interleaved.

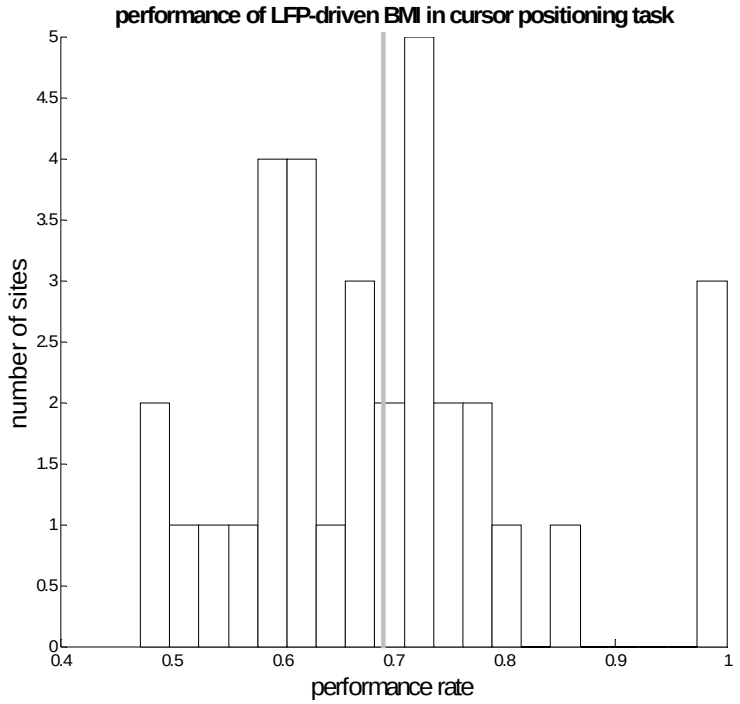

Figure 4-1 Histogram of LFP based BMI's performance.

4.3.2 Changes in LFP Characteristics with Cursor Use

In order to observe how the introduction of the cursor affected the neural activity we can emulate what the performance would have been using all the trials after the cursor was first introduced. Because the monkey does not know at the beginning of a trial whether

or not a cursor will appear, the offline analysis includes trials where the monkey completes reaches. Figure 2.1 displays the smoothed instantaneous performance of the classifier as a function of trial number for each site. The offline classifier uses the identical algorithm to the online classifier, cumulatively building the database, however the neural activity preceding not only cursor placements but reaches was classified according to predicted plan direction. The offline classifier performed similarly to the on-line classifier with a mean performance of 75.0%.

The vertical blue line indicates the time at which the cursor first appears. For most sites there is a trend towards improvement after the cursor is introduced. Two-thirds of the sites showed trends toward improved performance after the cursor was enabled, while one third of the sites became worse. On average the improvement in performance was 8.2%, while the average decrement in performance was 6.6%. We also performed the same analysis using a moving database (using only 15 trials prior to the decoded trial for each direction) with similar results.

Figure 4-3 is a histogram for the p-values for the MANOVA before and after the cursor is introduced. Note that although over 24 of the 33 sites showed significantly different mean activity for the 2 directions (MANOVA, $p<0.01$) both before and after the cursor utilization, only 75% of trials were predicted with accuracy.

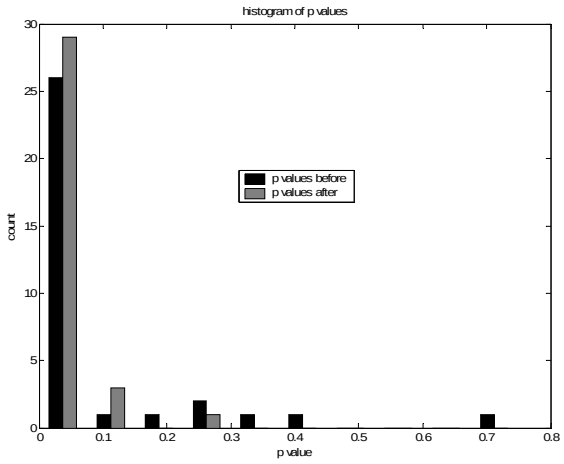

Figure 4-2 Histogram of MANOVA p-values for differences between reach directions before (black) and after (gray) BMI use.

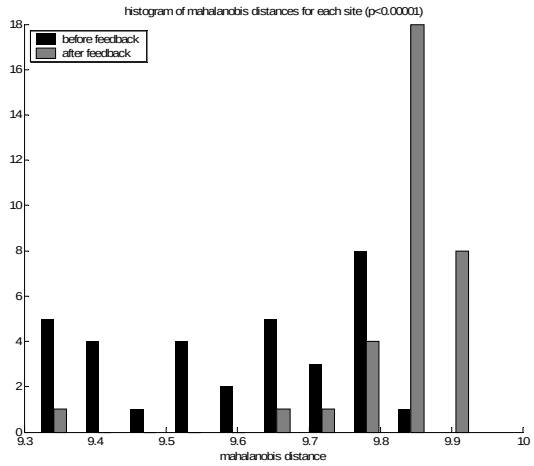

Figure 4-3 Histogram of mahalanobis distances between directions before (black) and after (grey) BMI use.

However, there was significant increase in the difference in neural activity between the two directions after cursor use. We used the mahalanobis distance as a metric to assess how introducing the cursor affected the difference between neural activity for the two reach directions. Larger mahalanobis distances indicate that the activity for the directions is dissimilar and more likely to be correctly classified.

Figure 4.4 shows a histogram of mahalanobis distances between the two directions before and after the cursor is introduced, only two sites show decreases in mahalanobis distance. The overall increase is significant (ANOVA, $p<0.000001$). It is clear that there is a significant change in the distances after the monkey starts using the cursor (kruskal-wallis, $p<0.00001$). Although the distribution of distances before cursor use is somewhat uniform, the distances after cursor use is limited primarily few bins, implying that there may be a ceiling effect that is reached rapidly over the course of each experiment.

4.3.3 Canonical Variate Analysis of LFP Features

In order to better understand the features of the LFPs that were used for classification we performed a canonical variate analysis on the entire dataset. The LFP signals were pre-processed for decoding to generate a 10 element feature vector of the average power contained in equal width frequency bands from 1 to 100 Hz. We constructed an optimal linear combination of these elements for each site. The maximum-likelihood method converges to the canonical optimization assuming the use of gaussian models for data distribution, as we used in this study. However, we have more explicit tools for observing the contributions of our features. The eigenvalues of the transformation matrix correspond to the ratio of between-group to within-group variance.

Figure 4-5 shows a typical example of data from the memory period as it looks after pre-processing. The diagonal shows the distribution of powers for each frequency band, and the lower triangle indicates the projection of each combination of features. Although some separation can be observed, there is clear overlap between the directions. (The artificial correlations in the 50-70 Hz bands are due to contamination by 60Hz.) By comparison, the separation between the optimally projected data in 4-6 finds easy distinction between directions. The red and blue lines plot the estimated source distributions of the corresponding data.

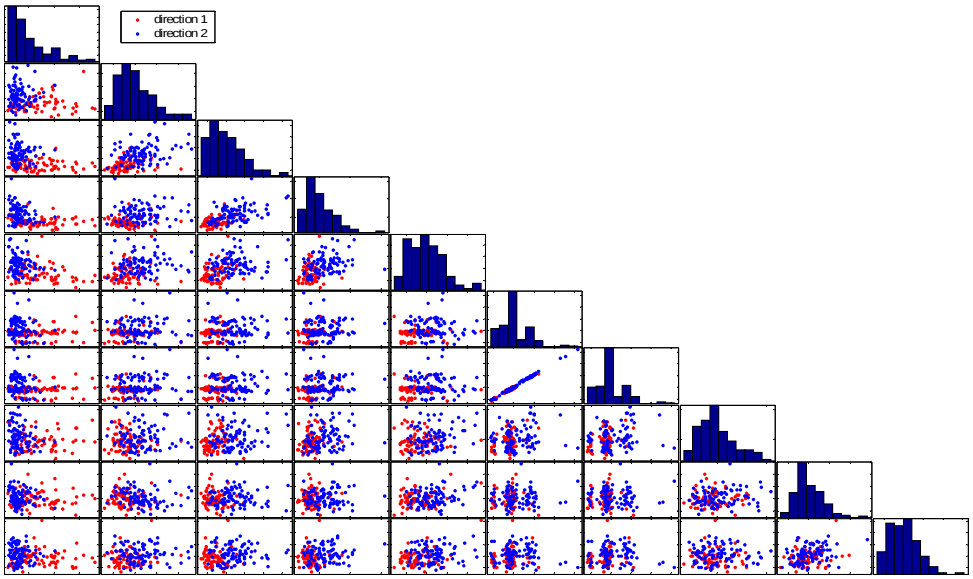

Figure 4-4 Pre-processed data for a single site. The diagonal contains the histogram of power for each frequency band (1-10, 11-20…91-100) during the memory period in normalized units, from top left to bottom right, respectively. Each scatter plot is the value of each of those features plotted against each other. The bottom left hand corner is the 1-10 Hz band plotted against the 91-100 Hz band. Each point represents the normalized power at the two compared frequencies for a single trial. Red trials are in direction 1, blue trials are in direction 2.

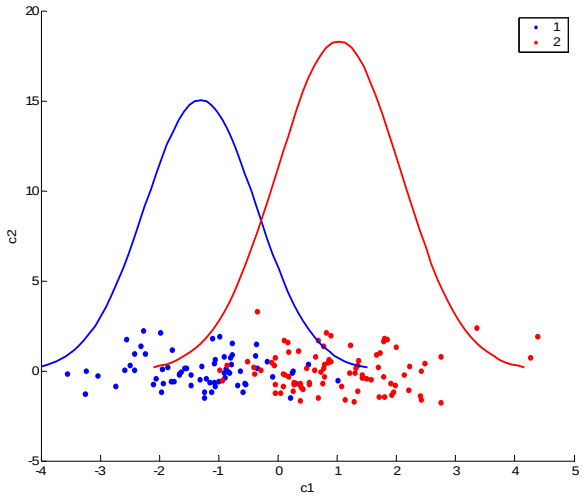

Figure 4-5 Data for direction 1 (blue) and direction 2 (red) projected optimally along first and second canonical variables (dots), with estimated probability distribution functions for the projection along the first canonical variable indicated by red and blue lines.

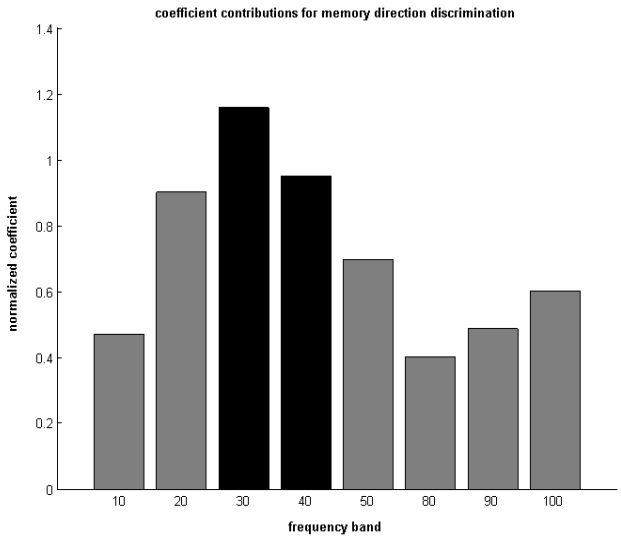

Figure 4-6 Normalized coefficient for the optimal linear transformation for each of the frequency bands. Bars in black represent those frequency bands that are significantly greater than the other bands.

In the two-group case, the only eigenvector of interest is the first. The values of the eigenvector correspond to the relative weighting of each of the original features. We can analyze the contributions of these features by observing the weights of the linear transformation. Figure 4-7 is a plot of the coefficients of the normalized optimal linear transformation; for each of the contributing frequency bands the coefficients of the 30-40 Hz bands are significantly greater (ANOVA, $p<0.001$) than the other bands (the bands contaminated by 60 Hz are excluded). The frequency band of the LFPs that carry information about the plan direction are primarily the 20-50 Hz bands.

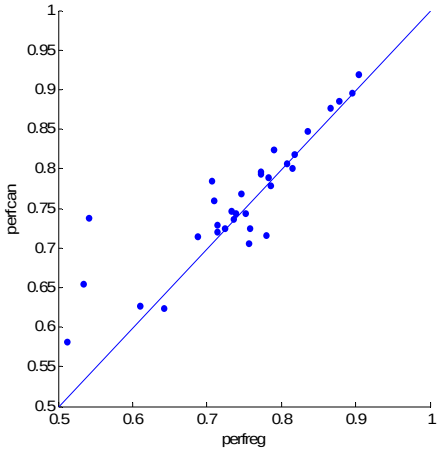

Figure 4-7 Improvement of performance using canonical variates vs. original data. All points above the diagonal represent performances that improve using canonical variates

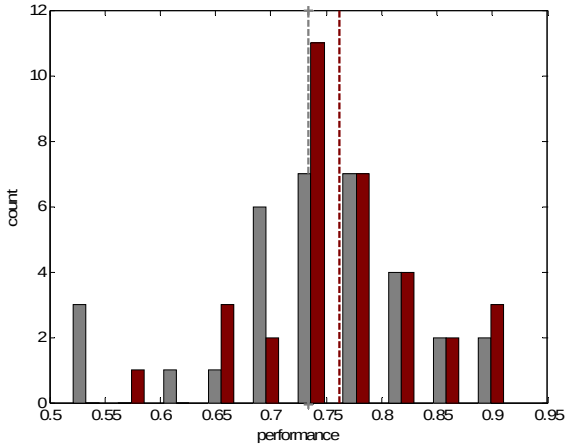

Figure 4-8 Histogram of performance of offline classifier on original data (gray) vs. canonical data (red).

Figure 4-8 compares the performance of the canonical variates as input to the classifier to the unweighted data; Figure 4-9 shows a histogram of the original data vs. the canonical data, showing a modest improvement.

We can repeat this analysis considering which frequency bands carry the most information about the location of the visual cue while it is illuminated vs. the memory period. The maximum-likelihood LFP based prediction of the cue location is comparable to the prediction of the reach plan. Figure 4-10 shows a histogram of this performance (76.6 ± 14.7). Figure 4-11 is a bar graph of the transformation coefficients for the cue for two directions during the cue period. The dominant frequency for cue direction discrimination is the 50 Hz band. This suggests that there is a type of 'multiplexing' in the LFPs, that different frequency bands carry different information at different points in the trial.

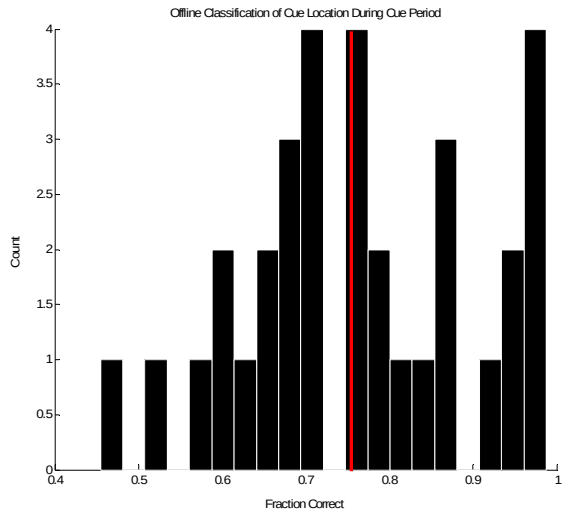

Figure 4-9 Histogram of prediction of direction from cue period data.

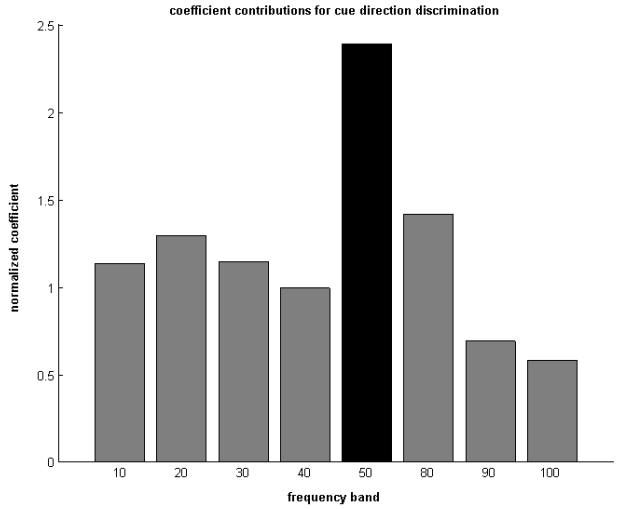

Figure 4-10 Relative weightings of each frequency band's contribution to the canonical variate used to determine the stimulus direction from data during the cue period across the population.

Another way of visualizing this is to perform the canonical analysis on data from the cue and the memory period, grouping by both direction and trial epoch. Figure 4-12 shows a scatter plot of the same site as in Figures 4-5 and 4-6, including both trial periods. The blue circles are from direction 1, the red circles from direction 2. Open circles demark the cue period, filled circles demark data from the planning period. In this combined situation, the best discrimination between plan directions is perpendicular to the horizontal axis, and the best discrimination between cue directions is perpendicular the vertical axis. The coefficients are correspondingly weighted and independent. In other words, the four clouds of points are not collinear, and the two orthogonal weighting schemes are necessary for discrimination. The axis of best directional discrimination independent of period is a line dividing the lower left and upper right quadrants of the data. The epoch itself (without considering direction) cannot be linearly discriminated in this projection.

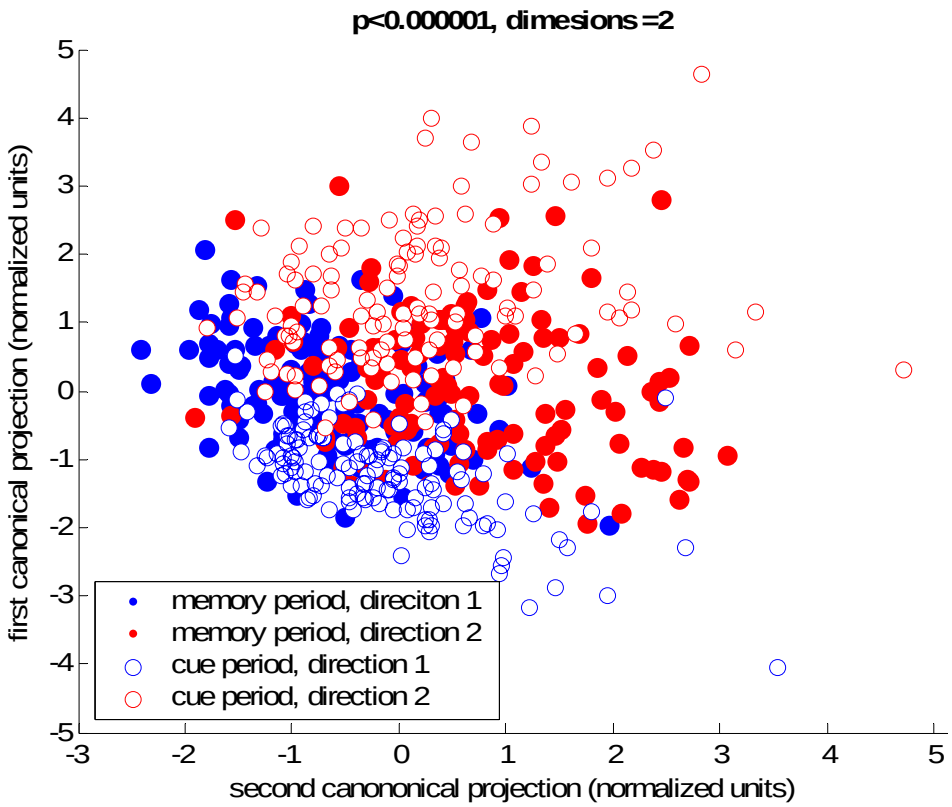

Figure 4-11 CVA projections onto first and second canonical variables as calculated including cue and memory period data for two directions.

We can also consider the evolution of the canonical variable over time and observe the difference between directions over time. Figures 4-13 and 4-14 show differences in canonical variables over the trial epoch for two different sites. The x-axis is time in trial, the y-axis the 95% confidence intervals of the means of the canonical variables for each direction. The means were calculated with 300ms windows overlapping by 150 ms. The vertical lines demark (in order) the cue being illuminated, the cue being extinguished, and the average start of the variably timed go signal (+/-200ms). Large differences indicate greater discriminability, smaller differences indicate weaker discriminability. The plot on the left is an example of a site whose peak discriminability is during the memory period,

whereas the plot on the right shows a site where there as much information about the trial direction during the cue as during the memory period.

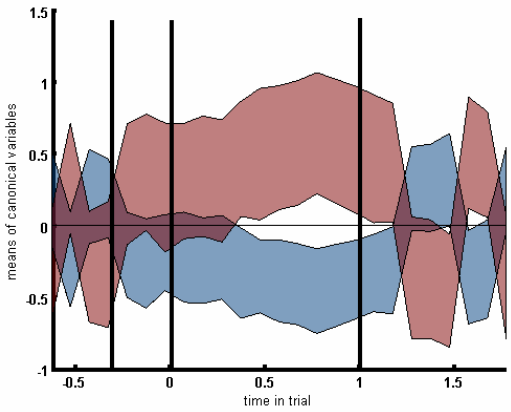

Figure 4-12

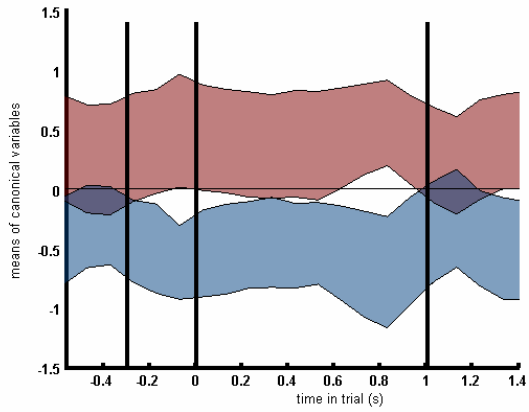

Figure 4-13

Figures 4-13 and 4-14. Evolution of canonical variables (normalized units) over the trial for two different sites. The different directions are indicated by the red and blue shaded areas. Black lines indicate cue illumination, the beginning of the memory period, and the average start time of the go signal (±0.2 s).

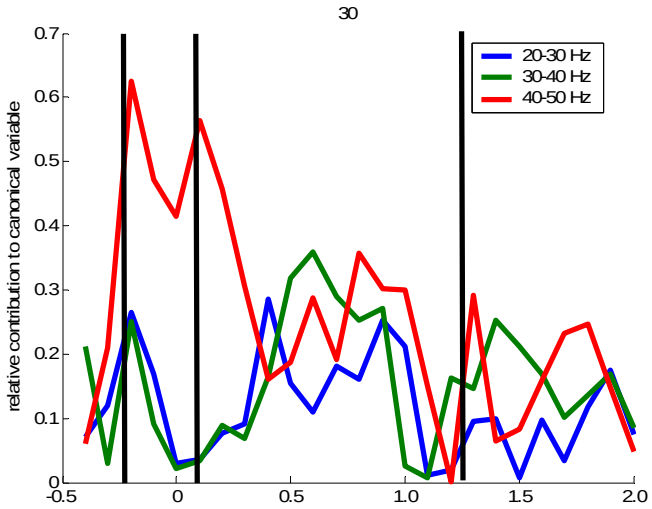

Figure 4-14 Evolution of contributions of different frequency bands to CVA over the course of the trial for a single site. Black lines demark the beginning of the cue period, the end of the cue period, and the (average) end of the memory period, from left to right. Vertical axis shows the relative contribution of each of the frequency bands shown, horizontal axis indicates time.

Overall, all of the 33 sites showed significant differences during the memory period, and 18 of those sites also showed significant differences during the cue period. 77.8% of these 18 sites showed higher discriminability during the cue period than the memory period, implying a separate 'population' of sites encode primarily visual (rather than plan) information.

Figure 4-15 indicates the progression of the contributions of the 20-30, 30-40, and 40-50 Hz to the continuously calculated canonical variable for the same site as figure 3.10, which showed significant discrimination for all task periods. During the cue period the 20 to 40 Hz contribute little to the discrimination, whereas the 50 Hz band is high. During the memory period, the relative contribution of the 50 Hz band drops while the contribution of the 20-40 Hz bands rise. This trend was repeated for 10 of the 17 sites where discrimination was significant for both the cue and the memory periods, though further investigation is necessary to clarify.

4.3.4 Saccades

For 12 of the 33 sites we also analyzed the saccadic activity. We found that the

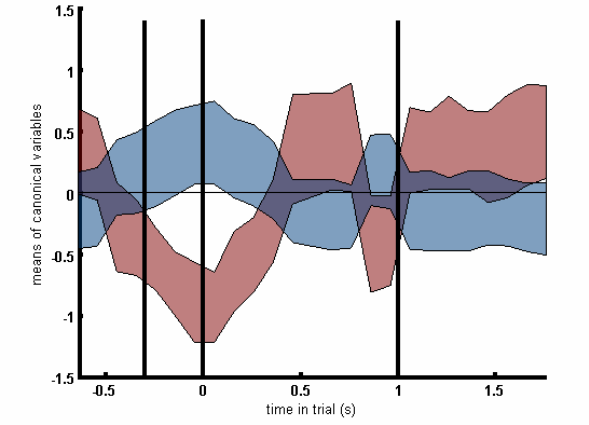

Figure 4-15 Evolution of means of canonical variates over the courser of a memory saccade trial. Black bars indicate the cue onset, cue extinction, and the (average) end of the memory period, from left to right. Red and blue lines indicate the 95% confidence intervals of the canonical variates as calculated for each 300ms sliding window throughout the trial.

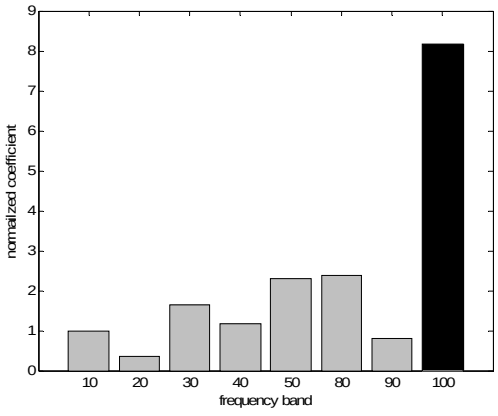

Figure 4-16 Relative contribution of selected frequency bands for 12 sites investigated for saccadic activity.

discrimination was poor between directions for most of the trial. The best discrimination was during the cue period. The 90-100 Hz band was the best for discrimination. Figure 4.1 shows the trial progression of the canonical means for saccades to the two selected *reach* directions at the same site as in figure 4-16 and 4-17. Figure 4-17 shows the average coefficients for the linear transformation producing the canonical variables across the population.

4.3.5 Effects of Selecting Sites without Action Potentials for LFP Data

The results we report here are from LFPs collected specifically from sites where no isolatable cells were present. The effect that this had on the discrimination between sites is displayed in figure 5.1. We compared the mahalanobis distance between the two directions measured in our dataset to that in a supplemental dataset. The supplemental dataset contained a sampling of electrode placements that included several sites with a variety of cellular proximities. We estimated the effect of spike presence by looking at the mahalanobis distance between the most separable reach directions as a function of spike amplitude on the electrode. We found that 10% of the variance in normalized mahalanobis distance can be accounted for by the spike waveform amplitude on the electrode. This indicates that these data show something of a lower limit for the performance of an LFP based BMI.

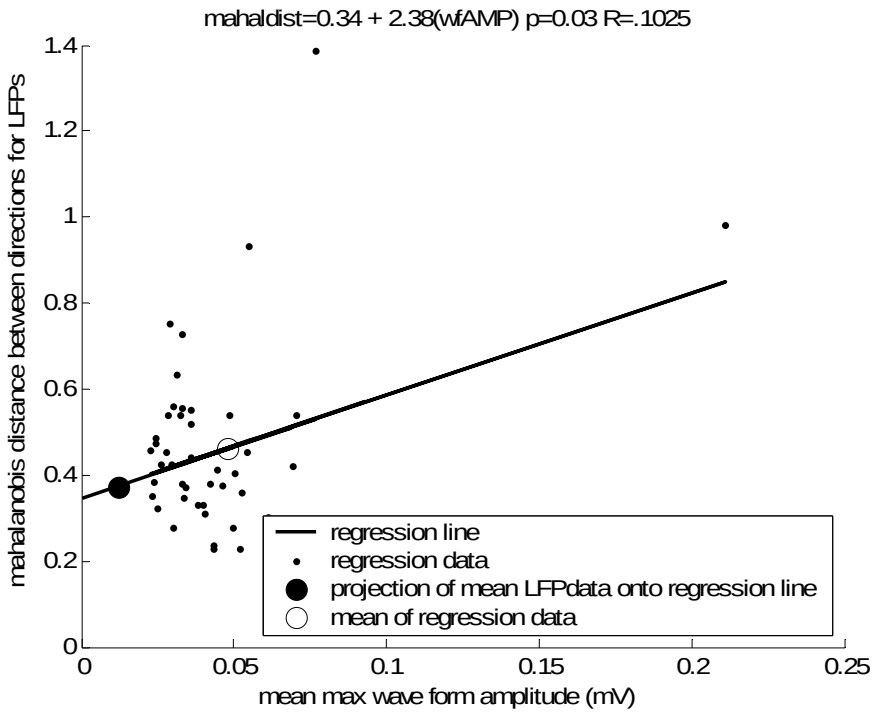

Figure 4-17 Effect of proximity to cell on discriminability. Regression plot for the mahalanobis distance between directions as a function of waveform amplitude on the same electrode. The regression line is in solid black, the data used for the regression are small black points, and the mean of the regression data is shown by the open circle. The projection of the LFP data used for the BMI onto the regression line is shown by the closed circle. Low values for this projection were expected, given that sites with no isolatable cells present were chosen.

4.3.6 Most Separable Directions

Figure 4-19 is a plot of the directions selected for feedback for each of the 33 experiments. Each line connects the selected directions for an experiment. Although most directions are 180° apart, many of the optimal separations came from sites that were separated by 135°.

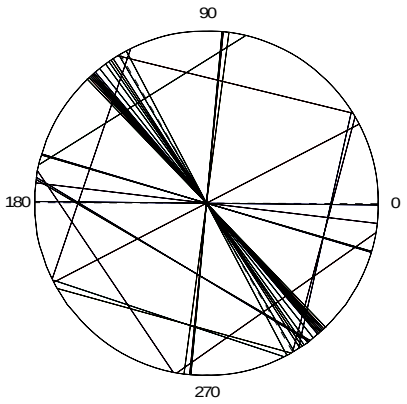

Figure 4-18 Black lines connect the two optimally separable directions selected out of eight for each site.

4.4 Discussion

The results of this study have implications for both the BMI community and for basic neuroscience. For the BMI community, we verify the utility of LFPs for BMI control. Many electrodes in implanted arrays do not have spikes present that can be well isolated, due either to electrode placement or deterioration of the local electrode environment with neural scarring. As has been seen here the LFPs can supplement the spiking data to great effect. This study presents the lower bound on LFP information (assuming that electrodes are placed in gray matter). The discrimination using LFPs from sites more proximal to cells should be equal or better than that presented here.

By further investigating the properties of the data we find that differing frequency bands contribute differently to discrimination of different types of information. The visual information is primarily contained in the 40-50 Hz frequency band, whereas the plan information is primarily contained in the 20-40 Hz bands. This is consistent with properties of the spike train spectra in the area (Buneo, Jarvis et al. 2003). This type of differentiation may prove useful in a BMI, when it is important not only to determine plan information (as was done in this experiment) but also the internal state of the patient.

Previous studies (Chapter 2/Appendix) show that the DC component of the LFPs in PRR are also useful for state determination. The lower frequency of the plan information is potentially indicative of specialized circuitry between parietal and premotor cortices. Additionally, the saccadic cue information is carried primarily in much higher frequencies, 90-100 Hz. Pesaran et al. showed that the gamma band frequencies in the LIP carried saccade target information (Pesaran, Pezaris et al. 2002). The power in these frequencies appears not to be maintained during the memory period in the the MIP, suggesting a gating of information from the LIP (or other upstream areas) to the MIP depending on the intention. This also indicates that the visual information is already interpreted and endowed with intentional significance at the point of its arrival in the MIP. The differentiation between the information-rich frequencies may root in afferents from upstream areas; the MIP receives inputs from V6, whereas the LIP is enervated strongly by MT/V5 (Rizzolatti 2003). These different input sources may affect the local circuitry differently, setting dynamics for the IPSPs (thought to dominate synchronization LFPs) in the local circuitry. Thus, the changes in synchronization reflected in LFPs may reflect entraining of downstream neurons to different frequencies of their inputs (Azouz and Gray 2000). Reach areas downstream of PRR also reflect spiking and LFP activity in the 16-50 Hz range, suggesting a preservation of this frequency signature across the reaching system (Murthy and Fetz 1996),(Donoghue, Sanes et al. 1998).

When we observe the evolution of directional discrimination throughout the trial, we see that in many cases the visual information is better distinguished than the plan information. Because LFPs are thought to be largely sourced from the local circuitry in a given area, the recordings where the visual information was more prominent may have been from areas whose local circuitry is strongly modulated by inputs from visual cortex, whereas the sites with strong plan discriminability may receive more local and feedback inputs from prefrontal areas. Given the trend exemplified in figure 3.10, it is possible that the known reciprocal connections between the extrastriate cortex and the MIP (Colby, Gattass et al. 1988) (Cavada and Goldman-Rakic 1993) (Felleman and Van Essen 1987; Boussaoud, Ungerleider et al. 1990) resonate (at 50 Hz) during and after the cue, while

those connections with frontal areas (Tanne-Gariepy, Rouiller et al. 2002) are engaged and maintain the plan information while there is no visual signal.

5 Prelimninary Evidence for Spike-Field Phase Model of Plasticity and Adaptation

This chapter presents a theory and some preliminary evidence for a link between spike-field phase, motivation, and learning. It is based on some basic concepts about the nature of LFP signals, spike timing dependent plasticity (STDP), and the role that reward and motivation may play in learning. It will not be attempted to fully review neuronal oscillation, plasticity, or reward literature, but to choose the most relevant elements of each for the purpose of presenting ways in which they might be causally related. Though extensive development of these ideas is necessary in future experiments, this preliminary model provides a structure for such experiments and makes some predictions given the limited data available thus far.

5.1 Background

Motivation and reward related modulations are found throughout the brain. One of the first studies identifying such signals was conducted in parietal cortex (Platt and Glimcher 1999), finding that LIP neurons activity was correlated with reward expectancy. Findings that increased reward improves discrimination in striatum (Hollerman, Tremblay et al. 1998), the prefrontal cortex (Matsumoto, Suzuki et al. 2003), and the parietal cortex (Musallam, Corneil et al. 2004) suggest that these types of reward based modulations could be linked with modifications that are known to occur with active discrimination practice (Zohary, Celebrini et al. 1994). Because such modifications do not accompany passive (unmotivated) stimulus exposure, the implication is that reward plays a key role in this type of discrimination improvement. Reward, in this context might be considered to be intrinsic or extrinsic. For example, in the increased BMI performance reported in Chapter 3, the reward might be considered intrinsic, insofar as less effort is necessary to operate the BMI than to execute a reach.

There is little evidence observing the effect of reward on LFPs in any of the cortical areas mentioned above. There has, however, been investigation of the role of spatial attention on spike-field coherency. Attention increased gamma frequency and reduced lower-

frequency synchronization of V4 neurons with the LFPs. Increasing the reward valance of a particular stimulus (or action) increases its attentional value, that is to say, that motivation, reward, attention, and learning are intrinsically linked (Kelley 2004). Though this attention study did not directly investigate any adaptive effects associated with attention and/or the synchronization observed, this type of modulation in LFP activity with attended (motivated) stimuli may also underlie the plasticity observed in similar tasks (Zohary, Celebrini et al. 1994).

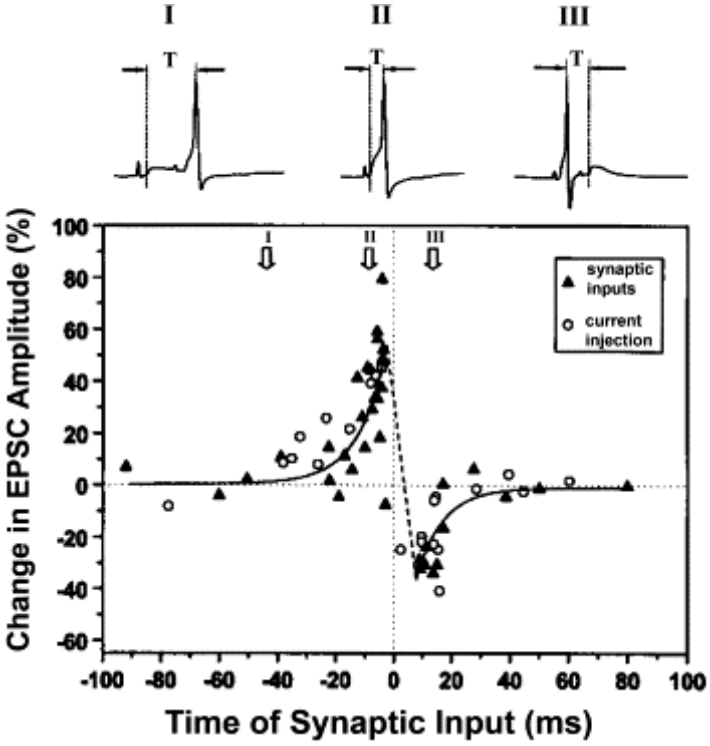

Figure 5-1 The critical window for spike timing-dependent plasticity of developing retinotectal synapses. The percentage change in the synaptic strength (EPSC amplitude) after repetitive retinal stimulation was plotted against the onset time of the retinal stimulation relative to the peak of the action potential initiated in the tectal cell. Data shown are for experiments in which spiking of the tectal neuron was initiated by either a suprathreshold input or a group of coactive inputs (filled circles) or by injection of a depolarizing current (open circles) in a tectal neuron (from (Zhang, Tao, et al. 1998)).

The concept of STDP has generated much interest for both theorists and experimentalists. For an excellent review of electrophysiology, proposed cellular mechanisms, and its relation to LTP and LTD as invoked by conventional methods, see (Dan and Poo 2004).

STDP has been characterized by an asymmetric window (Zhang, Tao et al. 1998), Figure 5-1.

The Hebbian "fire together, wire together" notion is refined by this data to indicate that causal relationships between neural firing result in synaptic potentiation, while acausal relationships result in synaptic depression. The amount of potentiation or depression is inversely related to the absolute value of the time lag between the pre- and postsynaptic neuronal spike times. A simple model of STDP among motion selective neurons predicted neuronal adaptation (in the cat visual cortex) and illusory perceptions in humans (Fu, Shen et al. 2004). Perhaps most relevant to this work, STDP of the excitatory connections in the rat hippocampus has been successfully modeled to mediate asymmetric expansion of neuronal place fields induced by repeated locomotion (Mehta, Quirk et al. 2000). The authors note that the phase precession (the advancement of the phase of the spike in 4-7Hz EEG oscillations) (O'Keefe and Recce 1993) is predicted by their model.

The connection between plasticity and spike-field coherency and phase is dependent on the three notions: First, that LFPs largely represent the IPSPs and EPSPs present at the electrode tip (Buzsaki and Draguhn 2004), close to the cell body, and are therefore most influential upon cell firing. Second, that STDP may be invoked based on inputs as inferred by the LFPs at the cell in question. Third, that synaptic plasticity that is invoked will affect the tuning of the cell. The consequential supposition is that if attention/reward/motivation alters (or accompanies an alteration in) the phase of the spikes and local field potentials, it will enable synaptic plasticity, and potentially a behavioral correlate, learning. In this case, the focus is on adaptations in cell tuning that might improve discrimination on a neuronal and/or psychophysical level.[4]

Below are three simple sketches of how tuning curve sharpness can be increased or decreased by synaptic potentiation and depression due to the phase of the cells' firing with its inputs (LFPs).

[4] In the case of the BMIs presented in the forgoing chapters, task performance (the typical psychophysical metric) and neuronal performance are identical.

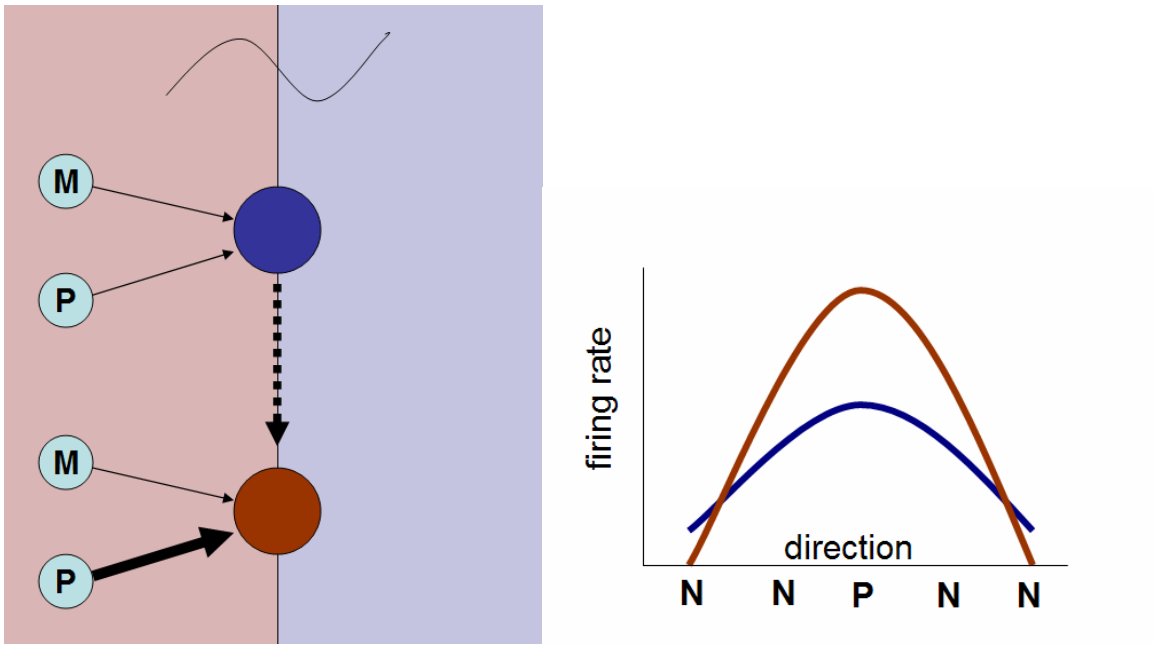

Figure 5-2 Causal firing can improve tuning by strengthening preferred direction inputs. The curve at the top of the image represents the spike triggered average of the LFP. The blue circles represent the cell of interest in the initial state of the 'network' and the red circles represent the cell of interest at the end state of the network. The M-inputs represent 'motivation' related input activity contributing to the LFP. P-units represent those inputs that are tuned similarly to the neuron. N-units represent inputs from units tuned differently from the neuron. The weights of the arrows represent the relative synaptic strengths that are incurred as a result of STDP. Units portrayed to the left or right figuratively fire before or after the cell of interest, respectively. **(A)** If the STA of the LFP in frequency ranges that are compatible with the time constant of the neuronal integration AND are tuned similarly to the neuron while motivation related inputs increase the likelihood of causal firing, input from the preferred direction should elicit stronger EPSCs from the cell of interest and **(B)** improve tuning; blue and red lines indicate hypothetical tuning curve before and after STDP respectively.

Figure 5-2 displays one scenario in which motivation related inputs might aid in sharpening the tuning of cells to improve discrimination. Similarly, tuning might be improved by depression of inputs from non-preferred directions, with a converse arrangement, Figure 5-3.

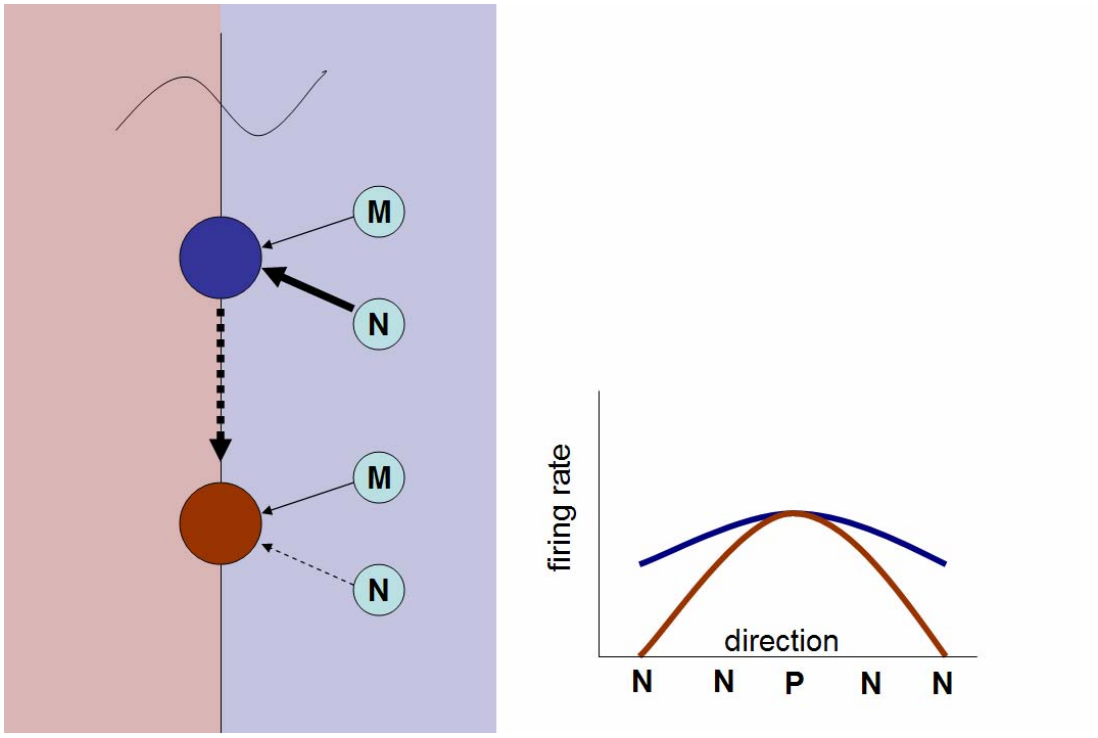

Figure 5-3 Acausal firing can improve tuning by depressing non-preferred direction inputs

If the inputs tuned similarly to the cell of interest tend to fire earlier and, a reward signal that causes a phase shift such that there is a causal relationship within the STDP integration window tuning will improve. The same tuning increase could also occur if there is an existing bias for inputs tuned in non-preferred directions to fire after the cell of interest within the STDP window. That is a motivation signal could shift any existing bias that a cell has for causal relationships between preferred inputs and acausal relationships with non-preferred inputs in order to increase tuning, assuming that it results in a positive phase. Given results that attention increases gamma-band coherence (Fries, Reynolds, et al. 2001), and it is gamma-band inputs (LFPs) that have been found to be tuned in PPC (Pesaran 2002), these predictions are compelling.

5.2 Preliminary Results

Limited data from recordings in MIP while reward conditions were varied were observed. The variable reward data was collected from an array of electrodes implanted in MIP for use with a BMI. Few cells were isolated with sufficient precision to assess spike-field

coherency, and neurophysiological and psychophysical performance metrics are inseparable. A monkey was trained to perform a memory reach task with cues for high (orange juice) and low (water) reward conditions given prior to the memory period. Data was evaluated during the memory period (1.2 to 1.8 s).

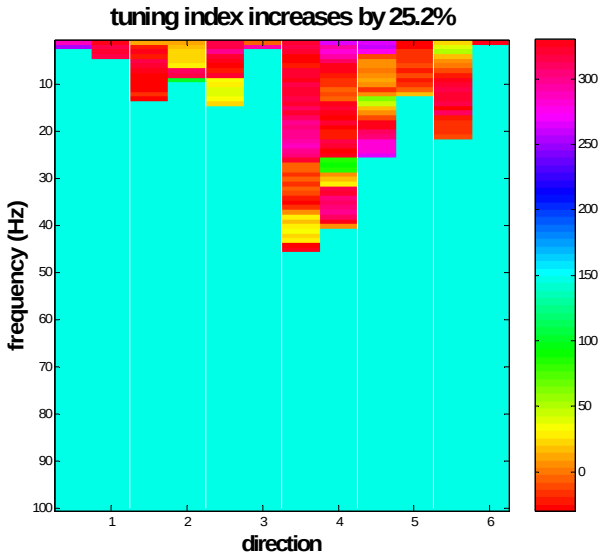

Figure 5-4. Phase and frequency for significantly tuned LFP frequencies in high (left bars of each column) and low (right bars of each column) reward cases. Frequency is represented on the vertical axis and direction is represented on the horizontal axis. Color indicates phase of the spiking to LFP.

Additionally, cells with low firing rates do not provide enough data to determine significant coherence. Here we present one of a few cells that were analyzable. Figure 5-4 shows the effect of reward level on LFPs for each direction. For most frequencies analyzed, there is no significant phase relationship. The preferred direction (direction 4) has more significant coherence estimates in higher frequencies due to more spikes available for analysis. Figure 5-5 shows the phase changes in low and high reward cases for trials in the preferred direction for 0-50 Hz for the same cell as above. The spike advances in phase relative to the lower frequencies, while retreating relative to the higher frequencies.

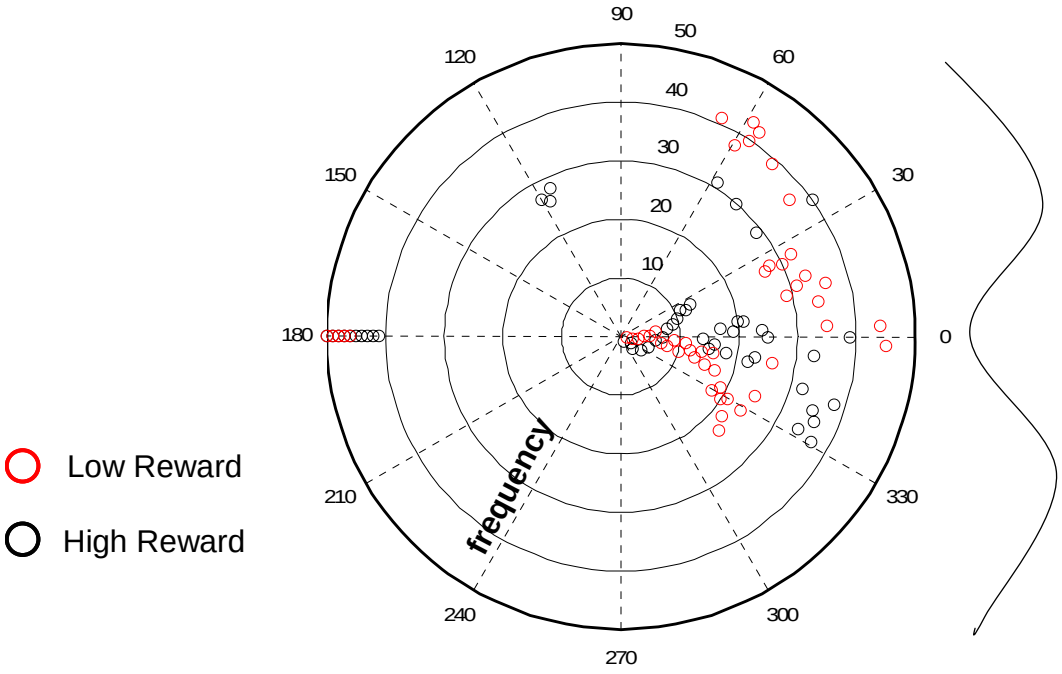

Figure 5-5 Phase diagram for high and low reward cases for trials in the preferred direction. Red points indicate the low reward condition while high frequencies indicate the high reward condition. Eccentricity represents increasing frequency (0-50Hz). Positive phases of spikes are present on spikes above the horizontal meridian, negative phase below the meridian. Zero phase is at the peak of the oscillation.

5.3 Implications

This proposal predicts some results regarding how reward affects the phase of spikes relative to inputs in order to induce STDP and in turn, discrimination. The phase of inputs that are tuned similarly to the cell (representing inputs from the preferred directions) will advance relative to spikes prior to tuning enhancement and phase retreat prior to tuning dampening. The phase of inputs that are tuned differently from the cell will advance prior to tuning damping, and retreat prior to tuning enhancement. "Motivation related" inputs should bring the cell closer to the firing threshold, making it easier for a cell to fire with presynaptic input, therefore, the phase of the motivation

signal with certain inputs can actually provide a rapid ‘gating’ mechanism changing the sensitivity of cells to different inputs.

These very preliminary results suggest that reward values effect the phase of spikes with respect to their inputs. Although the evidence is correlational, these changes in phase do accompany transitory tuning improvement between high and low reward conditions. The changes in phase are consistent with STDP, which ultimately could result in longer term plasticity. These preliminary results require extensive further research for confirmation, under conditions where spikes are well isolated, and psychophysical and neurphysiological discrimination metrics are separable. Better systems in which to test these ideas might be some of the cortical areas in which discrimination improvements have already been shown to exist, such as V1 and MT/MST. Mathematical and computer modeling of these predictions should be further developed concurrently to better guide experiments. Though outside the scope of this dissertation, further development could provide interesting insights and a better understanding of neural function and dynamics.

6 Conclusion

As stated in the introduction, producing brain machine interfaces raises both engineering and scientific challenges. The byproduct of the effort to solve these problems has been a better understanding of the underlying neuroscience. The studies described in the preceding chapters serve the intended purpose of demonstrating the feasibility and efficacy of parietal cortex as a locus for spike and LFP based BMIs as well as revealing features of how information is encoded in the brain. Here we discuss three ways in which this work has contributed to neuroscience and neuroengineering:

- Early offline studies described here (and in related studies) demonstrate that spike (and LFP) signals from PPC might be useful for BMI control.
- Follow-up studies implement PPC based BMIs using spikes in one case and LFPs in another.
- Collateral observations show adaptation effects with the use of PPC based BMIs, as well as elucidating other features of the neural signals.

These studies not only contribute to the advancement of BMI research, but systems neuroscience as a whole, and suggest several related future studies.

6.1 Offline Feasibility Studies

The offline study described in Chapter 2 was designed to determine whether spikes from PPC might be useful in the context of a BMI, and to introduce the concept of using a state machine approach to tracking internal state in the context of BMIs. Assuming ergodicity of serially collected multiple cells, this study demonstrates that applying simple transition rules predicting internal state to constrained task formats can generate accurate predictions from 8 possible targets and 3 different internal states with small numbers of cells (under 30). It also showed that LFPs were effective indicators of some of these state transitions. A related study (Scherberger et al., submitted) conducted similar offline predictions using LFPs recorded from sites where cells were present. The success of these studies laid the foundation for experiments in Chapters 3 and 4, where closed loop BMIs using spikes and LFPs respectively, were implemented.

6.2 Closed-Loop BMIs

Chapters 3 and 4 are both devoted to closed loop BMIs based on neural signals recorded from single electrodes placed in parietal cortex. The task differed from those in the feasibility studies in that only the movement direction was predicted and only two possible directions were tested. This was due to the use of a single electrode rather than multiple cells/sites. When BMIs using spikes were compared to BMI based LFPs from electrodes where spikes were *not* present, spike based BMIs were found to be 16% more accurate than LFP based BMIs in the same task. This weaker performance is mediated by the advantages of LFPs, such as the possibility that they contain information about internal state that is not present in spikes (as shown in Chapter 2), and the fact that they are easier to collect than spikes, using the same electrode technology. These studies were the first in their class to show that spikes/LFPs from PPC could be used effectively for BMIs.

Since the first study, work using multielectrode implants into PPC have shown that given the same task with as many as 6 possible target locations can be effectively decoded (Musallam, Corneil et al. 2004). This study is a logical extension of the original single electrode, two target condition. It not only shows that given multiple electrodes the target space can be expanded, but that electrode implants in PPC have longevity that is necessary for any human application of a neural prosthetic. Further studies show that PPC not only contains target information but also trajectory information (Mulliken 2004). The simple endogenous representation of target information in PPC (as opposed to what might be more complex information about movement dynamics represented in some downstream areas), taken together with these results and others outlined in the introduction, suggest a strong role for PPC in future development of BMIs.

6.3 Collateral Observations

6.3.1 Adaptation

Experiments understanding the role that PPC has in adaptation should not only shed light on adaptive processes in general but on the much debated role of PPC. Earlier studies have shown that LIP neurons will respond to sound after monkeys are trained to saccade to auditory targets (Linden, Grunewald et al. 1996). Other studies suggest that PPC plays an important role in prismatic adaptation (Clower, Hoffman et al. 1996). The results of the studies presented here imply still more ways in which plasticity of PPC might be engendered. Chapters 3 and 4 both indicate that there are adaptive changes that accompany use of BMIs for affected cells/sites. Roughly one-third of sites/cells show improvement in performance/tuning over the course of an experiment. This suggests that feedback of neural information can improve performance. Parallel observations that neural discrimination increases with reward (Musallam, Corneil et al. 2004) imply that in this context, the adaptation observed is in relation to an increased incentive for using the BMI. Results from human BMIs show promise that with the greater motivation and the communication ability of human subjects, this kind of improvement could be even further developed.

Chapter 5 shows preliminary evidence for the connection between motivation and synaptic modification in PPC.

6.3.2 LFP Informational Multiplexing

As part of the analysis in Chapters 2 and 4, it was found that LFPs in PPC contain information about trial features in addition to simple target location. Different frequency bands appear to contain information regarding internal state (i.e., a drop in 20Hz power prior to reaching), as well as differing specialization for effector type (eye vs. arms) and whether target information is visually available (as in the cue state) or simply retained during a planning phase. Though no study has been explicitly devoted to answering these questions, the results suggest that LFPs contain complex information, and different sites and frequency bands might be useful for interpreting information from upstream areas that is not integrated into spikes, but is significant nonetheless. This aspect of non-integration deserves separate research attention. Understanding why certain information

present in LFPs integrated into spiking information (and thus passed downstream) while other information remains influential at best is of interest. This unintegrated information may be a form of stochastic resonance, or may influence some of the learning mechanisms postulated in Chapter 5.

6.3.3 Adjusted Functional Definition of PPC Reach Areas

Tanaka and colleagues have found that bimodal neurons in parietal cortex expand their visual receptive fields to include areas of space that may be acted upon with the use of tools, but not within reach of the arm (Iriki A 1996). This adapataion was not present if the tool was simply held without a target present. They conclude that this is due to a subjective experience of the tool being assimilated to the hand. Visuo-motor transformations required for the use of a joystick similarly alter the mapping between natural movement domain and effector domain. Two studies that investigate the roles of motor cortex and dorsal premotor cortex suggest their differing roles in visuo-motor transformations (Shen and Alexander 1997); (Shen and Alexander 1997). Motor cortex neurons were found to be primarily tuned to movement trajectories. Over half of dorsal premotor neurons tested showed tuning to visual space during the presentation of a target that gave way to movement tuned activity during planning and reaching phases. Furthermore, the sensitivity of parietal areas to reward values indicates that the goal of the movement and its associated value are significant in PPC function (Musallam, Corneil et al. 2004); (Platt and Glimcher 1999). These studies suggest that unlike the motor cortex, cells in PPC code for targets *that can be manipulated by any means* in visual coordinates, a coordinate frame which is not altered by physical impairments or centered on subjective conditions.

BMI use, too, may be considered a more complex variety of tool use that requires 'visuo-motor' transformations that corresponds to minimal limb movement. The success of PPC BMIs, in particular trajectory based BMIs that are developed using visuo-motor transformations from a joystick to a cursor on a screen (Mulliken 2004), also suggest that the visual coordinate basis of parietal cortex may be very useful for BMIs which are designed for patients to act on space that would not be manually accessible[5]. It also may explain the contradicting and complex results that were found in the BMI associated plasticity of BMIs based primarily on motor cortical areas (Carmena, Lebedev et al.

2003); (Taylor, Tillery et al. 2002). That is, the type of visuo-motor transformations used in BMIs, tool extensions, and joysticks all have the shared property of transforming visual information to a motor instructional output. However, according to Tanaka and colleagues, once the manipulability of the targets is established, they gain visual representation in PPC, even if they are not part of the natural manual workspace. This might suggest a more inclusive definition/moniker for areas like PRR, which indicate effector independence in coding identified targets for immediate action. Future experiments should be conducted to refine this notion.

6.4 Future Experiments

The results of these experiments suggest future work, some of which is already underway both in the realm of neural prosthetics and basic neuroscience.

6.4.1 Basic Characterizations of Reach Related Areas in Parietal Cortex

The qualities of motor cortex have been extensively investigated under several reaching paradigms, but the relatively young area of research in the role of PPC in reach planning has been investigated primarily under paradigms originally designed in juxtaposition to established paradigms investigating area LIP. Several questions have yet to be answered about the basic features of neural activity related to reaching in PRR. Thorough investigations of direct reaches (without a delay period) with and without corresponding eye movement have not yet been conducted in PRR. Thus it is not clear whether simply training the unnatural task of delay reaches and/or fixation reaches might alter the endogenous representation in this area. There is little known about how depth information about reaching is encoded in PRR; this is of interest given the visual coordinate frame found in PRR, suggesting that vergence information may be encoded in PPC. Rajan Bhattacharyya has been undertaking experiments to investigate this issue. Although experiments investigating trajectory and obstacle avoidance are being conducted, another important area of study involves understanding the nature of how

[5] In unpublished observations, it was noted that comparable neuronal activity from an array implanted in MIP was present during purposeful movement of the leg as well as the arm.

continuous arm movement is represented in PRR in tasks that do not involve single targets. For example, [how] does PRR encode trajectory in a random smooth tracing task? Answering some of these basic questions may inform many of the more complex investigations currently being conducted.

6.4.2 Ventral-Dorsal Interactions

Chapter 1 alluded to a question of representation in PPC; that is, how do abstract forms, such as colors or shapes, come to be associated with particular movements and furthermore particular activity in PRR. This suggests an interaction between the ventral stream as a recognition engine and the dorsal stream as an action engine. It might be assumed from the LIP results that abstract forms without some significance as potential saccade or reach targets are not initially represented in PPC without a learning process. Chronic electrode arrays provide a way to investigate this issue. Though electrode arrays have been jointly placed in many areas of the dorsal stream, there have been few attempts to look at ventral-dorsal interactions. By identifying ventral elements that respond to particular features and then training associations of those features with eye or arm movements a clearer image of what the nature of these interactions are can be established. For example, the time course of learning particular associations and whether or not the associations are bidirectional should give rise to the locus of association between an abstract form and the action to which it corresponds.

6.4.3 BMI Related Plasticity

The several ways that these and other studies show that PPC may play a role in visuo-motor adaptation suggest further investigation of PPC under conditions such as prismatic adaptation, where a mismatch between visual space and motor space must be resolved. The chronic multiunit recording that is used in BMI research is ideal for answering these sorts of questions. Given the adaptive results from LIP, it would appear that explicit learning of new visuo-motor mappings would be represented clearly in PRR. However, this might not be the case in implicit learning paradigms, such as those involved in visuo-motor transformations of the entire space.

The contradictory adaptation results between the closed-loop studies by Taylor et al. and Carmena et al. beg the question of whether it is only the neurons that are used for feedback that adapt to improve the function of the BMI or if there are two states that are gated in the cortical areas involved, one for true reaches and another for BMI use. That is, if 'control' neurons not used for feedback were observed, would they show the same adaptive changes that those used for feedback do during BMI use, or would they retain the properties that are present during reaches? If this type of gating exists, what is its scope? Is it limited to neurons that are proximal functionally or spatially to those that do show adaptation, or does it affect larger cortical regions?

In the context of BMIs the notion that the so-called reaching areas of PRR might be redefined as workspace related areas suggests some experiments varying the nature of the space that can be manipulated by the BMI after its use has been trained. Such experiments may address similar issues to those of prismatic adaptation. Results indicate that if the space operated upon by the BMI were extended or altered, a PPC based BMI would adapt to include the additional representation more quickly than would motor cortex based BMIs.

6.4.4 Motivation, Reward, Learning, and LFPs

Chapter 5 presents a preliminary analysis and theory for the relationship between LFP and learning. A study designed to incorporate learning with motivational variables, such as reward valiance, with observation of learning and its neuronal correlates, is necessary to fully resolve and develop the notions presented. This could be done in a single electrode study that explores both the space of reward incentive and a learnable variable such as movement accuracy. If it can be observed that accuracy of reach movements refines cell tuning in PPC, as does motivation, different spaces of incentive and accuracy can be explored. The predictions of the preliminary research suggest that the phase of the

spikes in the 30-50 Hz range will change through tuning index increases and linking up information about STDP with motivation related adaptation.

Alternatively, experimental paradigms that have been shown to induce both learning and plasticity, such as those performed by Zohary and colleagues (Zohary, Shadlen et al. 1994), could be repeated while observing LFPs, perhaps from multiple electrodes, in order to generate better understanding of the population behavior and provide 'control' cells that do not show plasticity with the selected paradigm.

Mathematically formalizing the predictions and ideas set forth in Chapter 5 is also of great interest and will guide experiments.

Bibliography

Andersen, R. A., S. Musallam, et al. (2004). "Selecting the signals for a brain-machine interface." Curr Opin Neurobiol **14**(6): 720-6.

Andersen, R. A., L. H. Snyder, et al. (1997). "Multimodal representation of space in the posterior parietal cortex and its use in planning movements." Annu Rev Neurosci **20**: 303-30.

Azouz, R. and C. M. Gray (2000). "Dynamic spike threshold reveals a mechanism for synaptic coincidence detection in cortical neurons in vivo." Proc Natl Acad Sci U S A **97**(14): 8110-5.

Baizer, J. S., I. Kralj-Hans, et al. (1999). "Cerebellar lesions and prism adaptation in macaque monkeys." J Neurophysiol **81**(4): 1960-5.

Baker, J. T., J. P. Donoghue, et al. (1999). "Gaze direction modulates finger movement activation patterns in human cerebral cortex." J Neurosci **19**(22): 10044-52.

Barinaga, M. (1999). "Turning thoughts into actions." Science **286**(5441): 888-90.

Batista, A. P. and R. A. Andersen (2001). "The parietal reach region codes the next planned movement in a sequential reach task." J Neurophysiol **85**(2): 539-44.

Batista, A. P., C. A. Buneo, et al. (1999). "Reach plans in eye-centered coordinates." Science **285**: 257-260.

Batista, A. P., Buneo, C.A., Snyder, L.H. and Andersen, R.A. (1999). "Reach plans in eye-centered coordinates." Science **285**(5425): 257-260.

Birbaumer, N., A. Kubler, et al. (2000). "The thought translation device (TTD) for completely paralyzed patients." IEEE Trans Rehabil Eng **8**(2): 190-3.

Blatt, G., R. A. Andersen, et al. (1990). "Visual receptive field organization and cortico-cortical connections of area LIP in the macaque." Journal of Comparative Neurology **299**: 421-445.

Boussaoud, D. (1995). "Primate premotor cortex: modulation of preparatory neuronal activity by gaze angle." J Neurophysiol **73**(2): 886-90.

Boussaoud, D. and F. Bremmer (1999). "Gaze effects in the cerebral cortex: reference frames for space coding and action." Exp Brain Res **128**(1-2): 170-80.

Boussaoud, D., L. G. Ungerleider, et al. (1990). "Pathways for motion analysis: cortical connections of the medial superior temporal and fundus of the superior temporal visual areas in the macaque." J Comp Neurol **296**(3): 462-95.

Brown, E. N., L. M. Frank, et al. (1998). "A statistical paradigm for neural spike train decoding applied to position prediction from ensemble firing patterns of rat hippocampal place cells." J Neurosci **18**(18): 7411-25.

Buneo, C. A., M. R. Jarvis, et al. (2003). "Properties of spike train spectra in two parietal reach areas." Exp Brain Res **153**(2): 134-9.

Buneo, C. A., Jarvis, M. R., Batista, A. P., and Andersen, R. A. (2002). "Direct visuomotor transformations for reaching." Nature **416**: 632-636.

Buzsaki, G. and A. Draguhn (2004). "Neuronal oscillations in cortical networks." Science **304**(5679): 1926-9.

Caminiti, R., P. B. Johnson, et al. (1991). "Making arm movements within different parts of space: the premotor and motor cortical representation of a coordinate system for reaching to visual targets." J Neurosci **11**(5): 1182-97.

Caminiti, R., P. B. Johnson, et al. (1990). "Making arm movements within different parts of space: dynamic aspects in the primate motor cortex." J Neurosci **10**(7): 2039-58.

Carmena J. M., M. A. L., R. E. Crist, J. E. O'Doherty, D.M. Santucci, D. F. Dimitrov, P. G. Patil, C. S. Henriquez, M. A. L. Nicolelis (2003). "Learning to control a brain-machine interface for reaching and grasping by primates." PLoS **1**: 193-208.

Carmena, J. M., M. A. Lebedev, et al. (2003). "Learning to control a brain-machine interface for reaching and grasping by primates." PLoS Biol **1**(2): E42.

Cavada, C. and P. S. Goldman-Rakic (1993). "Multiple visual areas in the posterior parietal cortex of primates." Prog Brain Res **95**: 123-37.

Chapin, J. K., K. A. Moxon, et al. (1999). "Real-time control of a robot arm using simultaneously recorded neurons in the motor cortex." Nat Neurosci **2**(7): 664-70.

Cisek, P. and J. F. Kalaska (1999). "Neural correlates of multiple potential motor actions in primate premotor cortex." Society of Neuroscience Abstracts **25**: 381.

Clower, D. M., J. M. Hoffman, et al. (1996). "Role of posterior parietal cortex in the recalibration of visually guided reaching." Nature **383**(6601): 618-21.

Cohen, Y. E., Batista, A. P. and Andersen R. A. (2002). "Comparison of neural activity preceding reaches to auditory and visual stimuli in the parietal reach region." NeuroReport **13:**891-894.

Colby, C. L., R. Gattass, et al. (1988). "Topographical organization of cortical afferents to extrastriate visual area PO in the macaque: a dual tracer study." J Comp Neurol **269**(3): 392-413.

Connolly, J. D., M. A. Goodale, et al. (2000). "A comparison of frontoparietal fMRI activation during anti-saccades and anti-pointing." J Neurophysiol **84**(3): 1645-55.

Connolly, J. D., R. A. Andersen and M. A. Goodale (2003). "FMRI evidence for a 'parietal reach region' in the human brain." Exp Brain Res.

Crammond, D. J. and J. F. Kalaska (2000). "Prior information in motor and premotor cortex: activity during the delay period and effect on pre-movement activity." J Neurophysiol **84**(2): 986-1005.

Crist, R. E., W. Li, et al. (2001). "Learning to see: experience and attention in primary visual cortex." Nat Neurosci **4**(5): 519-25.

Cruz, V. T., B. Nunes, et al. (2003). "Cortical remapping in amputees and dysmelic patients: a functional MRI study." NeuroRehabilitation **18**(4): 299-305.

Dan, Y. and M. M. Poo (2004). "Spike timing-dependent plasticity of neural circuits." Neuron **44**(1): 23-30.

Donoghue, J. P., J. N. Sanes, et al. (1998). "Neural discharge and local field potential oscillations in primate motor cortex during voluntary movements." J Neurophysiol **79**(1): 159-73.

Farwell, L. A. and E. Donchin (1988). "Talking off the top of your head: toward a mental prosthesis utilizing event-related brain potentials." Electroencephalogr Clin Neurophysiol **70**(6): 510-23.

Felleman, D. J. and D. C. Van Essen (1987). "Receptive field properties of neurons in area V3 of macaque monkey extrastriate cortex." J Neurophysiol **57**(4): 889-920.

Fetz, E. E. (1969). "Operant conditioning of cortical unit activity." Science **163**(870): 955-8.

Fetz, E. E. (1999). "Real-time control of a robotic arm by neuronal ensembles." Nat Neurosci **2**(7): 583-4.

Fries, P., J. H. Reynolds, et al. (2001). "Modulation of oscillatory neuronal synchronization by selective visual attention." Science **291**(5508): 1560-3.

Fu, Y. X., Y. Shen, et al. (2004). "Asymmetry in visual cortical circuits underlying motion-induced perceptual mislocalization." J Neurosci **24**(9): 2165-71.

Galletti, C., P. Fattori, et al. (1997). "Arm movement-related neurons in the visual area V6A of the macaque superior parietal lobule." Eur J Neurosci **9**(2): 410-3.

Gauthier, G. M., J. M. Hofferer, et al. (1979). "Visual-motor adaptation: Quantitative demonstration in patients with posterior fossa involvement." Arch Neurol **36**(3): 155-60.

Georgopoulos, A. P., Kettner, R. E., and Schwartz, A. B. (1988). "Primate motor cortex and free arm movements to visual targets in 3-dimensional space. Coding of the direction of movement by a neuronal population." J Neurosci **8**(8): 2928-2937.

Ghez, C., J. Gordon, et al. (1995). "Impairments of reaching movements in patients without proprioception. II. Effects of visual information on accuracy." J Neurophysiol **73**(1): 361-72.

Gnadt, J. W. and R. A. Andersen (1988). "Memory related motor planning activity in posterior parietal cortex of macaque." Experimental Brain Research **70**: 216-220.

Goodale, M. A. and A. D. Milner (1992). "Separate visual pathways for perception and action." Trends in Neuroscience **15**: 20-25.

Goodale, M. A. and D. A. Westwood (2004). "An evolving view of duplex vision: separate but interacting cortical pathways for perception and action." Curr Opin Neurobiol **14**(2): 203-11.

Gray, C. M., P. Konig, et al. (1989). "Oscillatory responses in cat visual cortex exhibit inter-columnar synchronization which reflects global stimulus properties." Nature **338**(6213): 334-7.

Grunewald, A., J. F. Linden, et al. (1999). "Responses to auditory stimuli in macaque lateral intraparietal area. I. Effects of training." J Neurophysiol **82**(1): 330-342.

Hatsopoulos, N., J. Joshi, et al. (2004). "Decoding continuous and discrete motor behaviors using motor and premotor cortical ensembles." J Neurophysiol **92**(2): 1165-74.

Held, R. and A. Hein (1958). "Adaptation of disarranged hand-eye coordination contingent upon re-afferent stimulation." Perceptual and Motor Skills **8**: 87-90.

Helmholtz, H. (1867). Handbuch der Physiologischen Optik. Leipzig, Leopold Voss.

Hollerman, J. R., L. Tremblay, et al. (1998). "Influence of reward expectation on behavior-related neuronal activity in primate striatum." J Neurophysiol **80**(2): 947-63.

Hyvarinen, J. and A. Poranen (1974). "Function of the parietal associative area 7 as revealed from cellular discharges in alert monkeys." Brain **97**(4): 673-92.

Inoue, K., R. Kawashima, et al. (1997). "Activity in the parietal area during visuomotor learning with optical rotation." Neuroreport **8**(18): 3979-83.

Iriki A, T. M., Iwamura Y. (1996). "Coding of modified body schema during tool use by macaque postcentral neurones." Exp Brain Res **7**: 2325-2330.

Isaacs, R. E., D. J. Weber, et al. (2000). "Work toward real-time control of a cortical neural prothesis." IEEE Trans Rehabil Eng **8**: 196-198.

Johnson, P. B., S. Ferraina, et al. (1996). "Cortical networks for visual reaching: physiological and anatomical organization of frontal and parietal lobe arm regions." Cereb Cortex **6**(2): 102-19.

Karnath, H. O., P. Schenkel, et al. (1991). "Trunk orientation as the determining factor of the 'contralateral' deficit in the neglect syndrome and as the physical anchor of the internal representation of body orientation in space." Brain **114**(Pt 4): 1997-2014.

Kelley, A. E. (2004). "Memory and addiction: shared neural circuitry and molecular mechanisms." Neuron **44**(1): 161-79.

Kennedy, P. R. (1989). "The Cone Electrode - A Long-Term Electrode That Records from Neurites Grown onto Its Recording Surface." Journal of Neuroscience Methods **29**(3): 181-193.

Kennedy, P. R. and R. A. Bakay (1998). "Restoration of neural output from a paralyzed patient by a direct brain connection." Neuroreport **9**(8): 1707-11.

Kennedy, P. R., R. A. E. Bakay, et al. (2000). "Direct control of a computer from the human central nervous system." IEEE Transactions on Rehabilitation Engineering **8**(2): 198-202.

Kowler, E., J. van der Steen, et al. (1984). "Voluntary selection of the target for smooth eye movement in the presence of superimposed, full-field stationary and moving stimuli." Vision Res **24**(12): 1789-98.

Kurata, K. and E. Hoshi (1999). "Reacquisition deficits in prism adaptation after muscimol microinjection into the ventral premotor cortex of monkeys." J Neurophysiol **81**(4): 1927-38.

Lajoie, Y., J. Paillard, et al. (1992). "Mirror drawing in a deafferented patient and normal subjects: visuoproprioceptive conflict." Neurology **42**(5): 1104-6.

Linden, J. F., A. Grunewald, et al. (1996). "Auditory sensory responses in area LIP." Society of Neuroscience Abstracts **22**: 1062.

Marshall, J. C. and P. W. Halligan (1988). "Blindsight and insight in visuo-spatial neglect." Nature **336**(6201): 766-7.

Matsumoto, K., W. Suzuki, et al. (2003). "Neuronal correlates of goal-based motor selection in the prefrontal cortex." Science **301**(5630): 229-32.

Maunsell, J. H. and W. T. Newsome (1987). "Visual processing in monkey extrastriate cortex." Annu Rev Neurosci **10**: 363-401.

Maynard, E. M., N. G. Hatsopoulos, et al. (1999). "Neuronal interactions improve cortical population coding of movement direction." J Neurosci **19**(18): 8083-93.

McFarland, D. J., W. A. Sarnacki, et al. (2003). "Brain-computer interface (BCI) operation: optimizing information transfer rates." Biol Psychol **63**(3): 237-51.

Meeker, D., S. Cao, et al. (2001). Closed loop control of a neuroprosthetic. Society for Neuroscience, New Orleans, LA.

Mehta, M. R., M. C. Quirk, et al. (2000). "Experience-dependent asymmetric shape of hippocampal receptive fields." Neuron **25**(3): 707-15.

Milner, A. D. and M. A. Goodale (1993). "Visual pathways to perception and action." Prog Brain Res **95**: 317-37.

Mitz, A. R. and S. P. Wise (1987). "The somatotopic organization of the supplementary motor area: intracortical microstimulation mapping." J Neurosci **7**(4): 1010-21.

Mountcastle, V. B., J. C. Lynch, et al. (1975). "Posterior parietal association cortex of the monkey: command functions for operations within extrapersonal space." J Neurophysiol **38**(4): 871-908.

Moxon, K. A., N. M. Kalkhoran, et al. (2004). "Nanostructured surface modification of ceramic-based microelectrodes to enhance biocompatibility for a direct brain-machine interface." IEEE Trans Biomed Eng **51**(6): 881-9.

Mulliken, G. H., Musallam, S., Andersen, R.A. (2004). "Decoding trajectories in real time from posterior parietal cortex." Society of Neuroscience Abstracts.

Murata, A., V. Gallese, et al. (1996). "Parietal neurons related to memory-guided hand manipulation." J Neurophysiol **75**(5): 2180-6.

Murthy, V. N. and E. E. Fetz (1992). "Coherent 25- to 35-Hz oscillations in the sensorimotor cortex of awake behaving monkeys." Proc Natl Acad Sci U S A **89**(12): 5670-4.

Murthy, V. N. and E. E. Fetz (1996). "Synchronization of neurons during local field potential oscillations in sensorimotor cortex of awake monkeys." J Neurophysiol **76**(6): 3968-82.

Musallam, S., B. D. Corneil, et al. (2004). "Cognitive control signals for neural prosthetics." Science **305**(5681): 258-62.

Mushiake, H., Y. Tanatsugu, et al. (1997). "Neuronal activity in the ventral part of premotor cortex during target- reach movement is modulated by direction of gaze." J Neurophysiol **78**(1): 567-71.

Mussa-Ivaldi, S. (2000). "Real brains for real robots." Nature **408**(6810): 305-6.

Nicolelis, M. A., D. Dimitrov, et al. (2003). "Chronic, multisite, multielectrode recordings in macaque monkeys." Proc Natl Acad Sci U S A **100**(19): 11041-6.

Nordhausen, C. T., E. M. Maynard, et al. (1996). "Single unit recording capabilities of a 100 microelectrode array." Brain Res **726**(1-2): 129-40.

O'Keefe, J. and M. L. Recce (1993). "Phase relationship between hippocampal place units and the EEG theta rhythm." Hippocampus **3**(3): 317-30.

Paninski, L., M. R. Fellows, et al. (2004). "Spatiotemporal tuning of motor cortical neurons for hand position and velocity." J Neurophysiol **91**(1): 515-32.

Pesaran, B., J. Pezaris, et al. (2002). "Temporal structure in neuronal activity during working memory in Macaque parietal cortex." Nature Neuroscience **5**: 805-811.

Pesaran, B., Pezaris, J., Sahani, M., Mitra, P.M., and Andersen, R.A. (2002). "Temporal structure in neuronal activity during working memory in Macaque parietal cortex." Nature Neuroscience **5:**805-811.

Platt, M. L. and P. W. Glimcher (1999). "Neural correlates of decision variables in parietal cortex." Nature **400**(6741): 233-238.

Platt, M. L. and P. W. Glimcher (1999). "Neural correlates of decision variables in parietal cortex." Nature **400**(6741): 233-8.

Recanzone, G. H., M. M. Merzenich, et al. (1992). "Frequency discrimination training engaging a restricted skin surface results in an emergence of a cutaneous response zone in cortical area 3a." Journal of Neurophysiology **67**: 1057-1070.

Rizzolatti, G., Giacomo,Matelli,M,Massimo (2003). "Two different streams form the dorsal visual system: anatomy and functions." Experimental brain research. **153:** 146-157.

Rousche, P. J. and R. A. Normann (1998). "Chronic recording capability of the Utah Intracortical Electrode Array in cat sensory cortex." J Neurosci Methods **82**(1): 1-15.

Schoups, A., R. Vogels, et al. (2001). "Practising orientation identification improves orientation coding in V1 neurons." Nature **412**(6846): 549-53.

Schwartz, A. B. (2004). "Cortical neural prosthetics." Annu Rev Neurosci **27**: 487-507.

Schwartz, A. B., Kipke, D. R., Perepelkin, P. D. (1996). Cortical control for prosthetic devices. Proc. SPIE Int. Soc. Opt. Eng.

Serruya MD, H. N., Paninski L, Fellows M. R., Donoghue J. P. (2002). "Instant neural control of a movement signal." Nature **416**: 141-142.

Shen, L. and G. E. Alexander (1997). "Neural correlates of a spatial sensory-to-motor transformation in primary motor cortex." J Neurophysiol **77**(3): 1171-94.

Shen, L. and G. E. Alexander (1997). "Preferential representation of instructed target location versus limb trajectory in dorsal premotor area." J Neurophysiol **77**(3): 1195-212.

Shenoy, K. V., S. A. Kureshi, et al. (1999). Toward prosthetic systems controlled by parietal cortex. Society of Neuroscience Abstracts.

Shenoy, K. V., Meeker, D., Cao, S.Y., Kureshi, S.A., Pesaran, B., Buneo, C.A., Batista A.R., Mitra, P.P., Burdick, J.W., Andersen, R.A. (2003). "Neural prosthetic control signals from plan activity." Neuroreport **14**: 591-596.

Shoham, S., E. Halgren, et al. (2001). "Motor-cortical activity in tetraplegics." Nature **413**(6858): 793.

Smyrnis, N., M. Taira, et al. (1992). "Motor cortical activity in a memorized delay task." Exp Brain Res **92**(1): 139-51.

Snyder, L. H., A. P. Batista, et al. (1997). "Coding of intention in the posterior parietal cortex." Nature **386**(6621): 167-70.

Snyder, L. H., Batista, A.P., Andersen, R.A. (1997). "Coding of intention in the posterior parietal cortex." Nature **386**(6621): 167-170.

Tanne-Gariepy, J., E. M. Rouiller, et al. (2002). "Parietal inputs to dorsal versus ventral premotor areas in the macaque monkey: evidence for largely segregated visuomotor pathways." Exp Brain Res **145**(1): 91-103.

Taylor, D. M., S. I. H. Tillery, et al. (2002). "Direct cortical control of 3D neuroprosthetic devices." Science **296**(5574): 1829-1832.

Thompson, D. J. (1982). "Spectrum estimation and harmonic analysis." Proceedings of IEEE **70**: 1055-1996.

Toth, L. J. and J. A. Assad (2002). "Dynamic coding of behaviourally relevant stimuli in parietal cortex." Nature **415**(6868): 165-8.

Turner, J. A., J. S. Lee, et al. (2001). "Somatotopy of the motor cortex after long-term spinal cord injury or amputation." IEEE Trans Neural Syst Rehabil Eng **9**(2): 154-60.

Ungerleider, L. G. and M. Mishkin (1982). "Two cortical visual systems." Analysis of Visual Behavior. D. J. Ingle, M. A. Goodale and R. J. W. Mansfield. Cambridge, MA, MIT Press**:** 549-585.

Vaughan, T. M., J. R. Wolpaw, et al. (1996). "EEG-based communication: prospects and problems." IEEE Trans Rehabil Eng **4**(4): 425-30.

Vetter, P., S. J. Goodbody, et al. (1999). "Evidence for an eye-centered spherical representation of the visuomotor map." J Neurophysiol **81**(2): 935-9.

Wallman, J. and A. F. Fuchs (1998). "Saccadic gain modification: visual error drives motor adaptation." J Neurophysiol **80**(5): 2405-16.

Wannier, T. M., M. A. Maier, et al. (1989). "Responses of motor cortex neurons to visual stimulation in the alert monkey." Neurosci Lett **98**(1): 63-8.

Weiner, M. J., M. Hallett, et al. (1983). "Adaptation to lateral displacement of vision in patients with lesions of the central nervous system." Neurology **33**(6): 766-72.

Welch, R. B., B. Bridgeman, et al. (1993). "Alternating prism exposure causes dual adaptation and generalization to a novel displacement." Percept Psychophys **54**(2): 195-204.

Wessberg, J., C. R. Stambaugh, et al. (2000). "Real-time prediction of hand trajectory by ensembles of cortical neurons in primates." Nature **408**(6810): 361-365.

Williams, J. C., R. L. Rennaker, et al. (1999). "Long-term neural recording characteristics of wire microelectrode arrays implanted in cerebral cortex." Brain Res Brain Res Protoc **4**(3): 303-13.

Wise, S. P., S. L. Moody, et al. (1998). "Changes in motor cortical activity during visuomotor adaptation." Exp Brain Res **121**(3): 285-99.

Wolpaw, J. R., N. Birbaumer, et al. (2000). "Brain-computer interface technology: a review of the first international meeting." IEEE Trans Rehabil Eng **8**(2): 164-73.

Wolpaw, J. R. and D. J. McFarland (2004). "Control of a two-dimensional movement signal by a noninvasive brain-computer interface in humans." Proc Natl Acad Sci U S A **101**(51): 17849-54.

Wolpert D.M., a. Z. G. (2000). "Computational principles of movement neuroscience." Nat Neurosci **3**: 1212-1217.

Zhang, K., I. Ginzburg, et al. (1998). "Interpreting neuronal population activity by reconstruction: unified framework with application to hippocampal place cells." J Neurophysiol **79**(2): 1017-44.

Zhang, L. I., H. W. Tao, et al. (1998). "A critical window for cooperation and competition among developing retinotectal synapses." Nature **395**(6697): 37-44.

Zohary, E., S. Celebrini, et al. (1994). "Neuronal plasticity that underlies improvement in perceptual performance." Science **263**(5151): 1289-92.

Zohary, E., P. Hillman, et al. (1990). "Time course of perceptual discrimination and single neuron reliability." Biol Cybern **62**(6): 475-86.

Zohary, E., M. N. Shadlen, et al. (1994). "Correlated neuronal discharge rate and its implications for psychophysical performance." Nature **370**(6485): 140-3.