BOSTON UNIVERSITY

SCHOOL OF MEDICINE

Dissertation

LEARNING AND ADAPTATION IN BRAIN MACHINE INTERFACES

by

SPENCER BRADLEY TORENE

B.A., University of Maryland, 2004

Submitted in partial fulfillment of the

requirements for the degree of

Doctor of Philosophy

2017

Approved by

First Reader _____

Jason T. Ritt, Ph.D.
Assistant Professor of Biomedical Engineering

Second Reader _____

Frank H. Guenther, Ph.D.
Professor of Speech, Language & Hearing Sciences and Biomedical Engineering

Third Reader _____

Timothy J. Gardner, Ph.D.
Assistant Professor of Biology and Biomedical Engineering

Humanty Doomanty perched on a wall,

Humanty Doomanty had a great fall,

All of the circuits and all of the mem

Took humanity's place from them.

— anonymous meatbag

DEDICATION

I would like to dedicate this work to my patient spouse Becca, my wonderful children Elliot and Simon, and my dog Sam.

ACKNOWLEDGMENTS

While I owe many thanks to many people, those singled out here deserve more than mere gratitude. Unfortunately for them, gratitude is the currency at present. First, thanks to my advisors, Jason and Frank, for guiding me through the third stage of learning how to think. I thank them for their scientific insight and wisdom, which has proven and will continue to prove invaluable. Thanks to Frank for bringing perspective and reason, and to Jason for consistently steering me towards better ways of answering questions. Thanks to my wife Rebecca for inspiring me to go back to graduate school, and who took Elliot to the trampoline park and numerous other locations while 8 months pregnant with Simon so I could write this dissertation. I would also like to thank Kayle Sawyer and EmilyKate McDonough, who told me to just use everything now and edit later, actually starting me on this daunting document. Thanks to my committee, Tim Gardner and Mark Kramer for bearing with me, and Chris Moore for providing extremely useful and encouraging comments on Chapter 2. Finally, many thanks to Shelley Russek and Sandi Grasso for having my back. They are the best.

LEARNING AND ADAPTATION IN

BRAIN MACHINE INTERFACES

SPENCER BRADLEY TORENE

Boston University School of Medicine, 2017

Major Professor: Jason T. Ritt, Ph.D., Assistant Professor of Biomedical Engineering

ABSTRACT

Balancing subject learning and decoder adaptation is central to increasing brain machine interface (**BMI**) performance. We addressed these complementary aspects in two studies: (1) a learning study, in which mice modulated "beta" band activity to control a 1D auditory cursor, and (2) an adaptive decoding study, in which a simple recurrent artificial neural network (**RNN**) decoded intended saccade targets of monkeys.

In the learning study, three mice successfully increased beta band power following trial initiations, and specifically increased beta burst durations from 157 ms to 182 ms, likely contributing to performance. Though the task did not explicitly require specific movements, all three mice appeared to modulate beta activity via active motor control and had consistent vibrissal motor cortex multiunit activity and local field potential relationships with contralateral whisker pad electromyograms. The increased burst durations may therefore by a direct result of increased motor activity. These findings suggest that only a subset of beta rhythm phenomenology can be volitionally modulated (e.g. the tonic "hold" beta), therefore limiting the possible set of successful beta neuromodulation strategies.

In the adaptive decoding study, RNNs decoded delay period activity in oculomotor and working memory regions while monkeys performed a delayed saccade task. Adaptive decoding sessions began with brain-controlled trials using pre-trained RNN models, in contrast to static decoding sessions in which 300-500 initial eye-controlled training trials were performed. Closed loop RNN decoding performance was lower than predicted by offline simulations. More consistent delay period activity and saccade paths across trials were associated with higher decoding performance. Despite the advantage of consistency, one monkey's delay period activity patterns changed over the first week of adaptive decoding, and the other monkey's saccades were more erratic during adaptive decoding than during static decoding sessions. It is possible that the altered session paradigm eliminating eye-controlled training trials led to either frustration or exploratory learning, causing the neural and behavioral changes.

Considering neural control and decoder adaptation of BMIs in these studies, future work should improve the "two-learner" subject-decoder system by better modeling the interaction between underlying brain states (and possibly their modulation) and the neural signatures representing desired outcomes.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

LIST OF ABBREVIATIONS

AAC ..... Amplitude-Amplitude Coupling

BI ..... Behavior Index

BMI ..... Brain Machine Interface

BPR ..... Beta Power Ratio

CFC ..... Cross Frequency Coupling

dlPFC ..... Dorsolateral Prefrontal Cortex

EEG ..... Electroencephalogram

EMG ..... Electromyogram

ERD ..... Event Related Desynchronization

ERS ..... Event Related Synchronization

FA ..... False Alarm

FEF ..... Frontal Eye Field

FWHM ..... Full-Width Half-Max

ISI ..... Inter-Spike Interval

ITI ..... Inter-Trial Interval

LDA ..... Linear Discriminant Analysis

LFP ..... Local Field Potential

MUA ..... Multi-unit Activity

PAC ..... Phase-Amplitude Coupling

RMSE ..... Root Mean Square Error

RNN ..... Simple Recurrent Neural Network

SEF ..... Supplementary Eye Field

SMR ..... Sensorimotor Rhythm

TP ..... True Positive

vM1 ..... Vibrissal Motor Cortex

CHAPTER ONE: Learning and Adaptation in Brain Machine Interfaces

Introduction

Brain-machine interfaces (**BMI**s) are a proposed technology to provide mobility and independence to paralyzed and "locked-in" patients. However, a major limitation to the viability of BMIs as a medical treatment is low performance. Two fundamental and complementary components of BMI performance are the subject's ability to learn to use the decoder, and the decoder's adaptation to changes in the subject's neural activity. We investigated both aspects in this dissertation and report on the following: (1) a BMI learning study in which mice controlled an auditory neurofeedback cursor, and (2) an adaptive BMI study in which a decoder was updated online as monkeys performed a delayed saccade task. Through this work we may better understand the distinct, but connected, roles of subject learning and decoder adaptation to create recommendations for higher BMI performance.

The learning study focused on an important prerequisite of motor control: the decision to move in the first place. Patients may need "asynchronous" BMIs that determine when subjects want to be in rest or motor execution states, in contrast to common "synchronous" BMIs, which assume subjects intend to move during experimenter defined control periods. We established whether the beta sensorimotor rhythm (**SMR**), a classic neural signature of motor preparation, might be used as an

asynchronous control signal by training mice to increase beta SMR power in response to auditory neurofeedback.

The adaptive study arose from an ongoing saccade BMI study (Brincat et al., 2013b). A BMI that detects the direction of intended saccades could be extremely useful to locked-in patients, who generally have difficulty moving their eyes. Previously employed static decoders required hundreds of training trials each day, placing a large burden on the subjects. As a possible improvement, we implemented an adaptive recursive artificial neural network (**RNN**) that could reduce the need for training, and improve decoding performance over trials. Although preliminary offline simulations suggested good performance was possible, the closed loop RNN was lower performing than both the preliminary offline results and the closed loop static LDA decoder. We assessed whether changes in monkey behavior or internal cognitive factors drove poor performance.

Subject Learning of BMI Control (Chapter Two)

Non-adaptive BMI studies, in which the decoder is fixed, allow subjects to achieve high levels of BMI skill (Ganguly and Carmena, 2009). Important work has looked at the ability of individual neurons to modulate activity and tuning properties in BMI tasks (Fetz, 1969; Jarosiewicz et al., 2008; Ganguly et al., 2011; Chase et al., 2012; Koralek et al., 2012; Clancy et al., 2014). However important these studies are, relying on unit activity to decode motor intent may not be suitable over long term timescales, due to gliosis or electrode degradation (Chestek et al., 2011; Prasad et al., 2012). The viability

of using LFPs—rather than unit activity—as a decoding signal should be investigated, due to the robustness of LFP decoding over the life of the electrode (Chao et al., 2010; Flint et al., 2013). Furthermore, LFPs have been shown to contain more information than unit activity about behavioral state when fewer than 20 electrodes are available (Scherberger et al., 2005; Bansal et al., 2012; Flint et al., 2012). Many studies have looked at the capability of subjects to modulate various LFP bands or SMRs (Wyrwicka and Sterman, 1968; Wolpaw and McFarland, 2004; Bai et al., 2008; Hwang and Andersen, 2009; Pfurtscheller et al., 2010; Engelhard et al., 2013; Rouse et al., 2013; Wander et al., 2013), but fewer have been devoted to understanding the learning process through which volitional LFP modulation takes place, including relevant behavioral observations (Rouse et al., 2013; Orsborn et al., 2014; So et al., 2014).

The ability of subjects to learn how to use decoders is a vital component of successful BMI systems, and our understanding of this process with LFPs is no less important. The motor cortex beta sensorimotor rhythm (**SMR**) is a good candidate for volitional modulation in motor BMIs, as a classic neural signature of sensorimotor activity (Sanes and Donoghue, 1993; Murthy and Fetz, 1996) that appears prominently in the entire cortical-basal ganglia loop (Cassidy et al., 2002; Kühn et al., 2004). The beta SMR has been successfully modulated in electroencephalography BMIs (Wolpaw and McFarland, 2004; Naros and Gharabaghi, 2015), but EEG cannot provide low level detail of the beta SMR modulation learning process. Up-regulation of the beta SMR is a more robust phenomenon than the beta ERD during both actual and imagined movement (Leeb et al., 2007; Bai et al., 2008; Pfurtscheller and Solis-Escalante, 2009; Pfurtscheller et al.,

2010; Solis-Escalante et al., 2010) and does not require overt movement (Pfurtscheller et al., 2005). Furthermore, down-regulation of neural activity appears to be harder than up-regulation, in general (Rouse et al., 2013; Clancy et al., 2014).

We trained mice to increase beta band power in the vibrissal motor cortex to receive water reward. Mice increased the rate of water rewards received per session through successful neuromodulation, and not through a successful behavioral strategy. It appears that the neuromodulation may have been accomplished through the extension of low frequency (alpha and beta) burst durations and was accompanied by particular EMG signatures.

BMI Decoder Adaptation (Chapter Three)

Adaptive BMI studies have shown that performance can improve with adaptation (Orsborn et al., 2014), and trajectory precision drops as soon as adaptation ceases and the decoder is fixed (Orsborn et al., 2012; Flint et al., 2013). Fixed decoders might allow performance to stagnate (Bishop et al., 2014), unless only a small portion of the most stable neurons are utilized (Ganguly and Carmena, 2009). Unfortunately, the recording quality of units do not always remain stable or consistent (Chestek et al., 2011; Prasad et al., 2012) and the information and performance that can be extracted from neural recordings decreases with reduced numbers of neurons (Wessberg et al., 2000; Carmena et al., 2003; Ganguly and Carmena, 2009). Performance could be better if adaptive decoding methods are used that minimize instability in neural recordings over time.

Therefore, increasing the cost-benefit ratio of intracortical implants will likely require decoder adaptation.

Even when decoders have good offline performance, their parameters change drastically after subject learning (Orsborn et al., 2012). Studies that recalibrate their decoders on a daily basis show high variability in performance (Taylor et al., 2002; Carmena et al., 2003; Gilja et al., 2012); it would be better to create decoders that had more consistent, high initial performance. Adaptation must be performed with caution, as even small changes in parameter weights can be highly detrimental to performance (Ganguly and Carmena, 2009). The timescale of adaptation has been researched before, from only one batch update (Gilja et al., 2012) to updates at every time step (Shpigelman et al., 2008; Vidaurre et al., 2011) to intermediate time scale updates (Orsborn et al., 2012). Whether it is desirable to use frequent or infrequent updating may depend on how much learning you can expect subjects to achieve, or on the quality of the implants. There is some debate as to whether effective learning can occur with frequent parameter updates (Li et al., 2011; Orsborn et al., 2014), although it appears that fixed or more slowly adapting decoders allow more refined movements to develop (Danziger et al., 2009), as well as biomimetic decoders as a starting point (Sadtler et al., 2014). Learning tends to occur as demand dictates—that is, units that have improper assignments of preferred direction tend to shift those preferred directions, and units that are weighted more heavily increase their modulation depth (Orsborn et al., 2014). While there is some debate as to what the best adaptation paradigm is (DiGiovanna et al., 2009; Li et al., 2011; Gilja et al., 2012; Orsborn et al., 2012; Dangi et al., 2014; Marsh et al., 2015), it

appears clear that gradual modifications of biomimetic decoders are most intuitively learnable by subjects (Danziger et al., 2009; Sadtler et al., 2014). To this end, we trained a simple recurrent artificial neural network (**RNN**) with several generalization techniques on previously recorded data as a starting point for further adaptation of a delayed saccade BMI.

Performance of the closed loop adaptive RNN was worse than both the offline RNN adaptation and a closed loop static LDA decoder. We suspect the poor adaptive performance during the delayed saccade task was due to changes in the monkeys' strategies and motivation at the loss of initial eye control trials, and is not consistent with poor decoder adaptation parameters.

Common Themes

There are two common themes in these studies: (1) the use of LFPs as decoding features, and (2) the study of motor intention and behavior.

LFPs at the opposite ends of the neural spectrum were used as features to decode different types of motor intention and activities on different timescales. In the learning study we looked at how low frequency LFPs (13-30 Hz) can be volitionally modulated over several seconds to inform decoders of an overall, broad subject motor state. In the adaptation study we looked at how high frequency LFPs (80-500 Hz) within a brief, time-limited delay period can be decoded into dynamic motor activity. Furthermore, in the adaptive work, we briefly attempted to extract additional neural information by modeling the relationship between the low and high frequency LFP bands. The diversity of

processes for which various LFP frequency bands may inform decoders (Fontolan et al., 2014; Bundy et al., 2016; Gao, 2016; Michalareas et al., 2016) is illustrated in these works as potential surrogates of missing spiking activity.

Another common theme is the study of motor intention. In the learning study we used beta SMR frequencies as a "switch" for the intent to move and found that there was a direct connection between motor cortical activity and EMG activity. This connection was possibly a strategy or an outcome of learning the specific task requirement to generate high beta power. In the adaptation study we decoded delay period working memory to guess the directional intention of upcoming eye movements, the success of which was related to the stereotypy of post-delay saccades. In both studies we report on task-related motor behavior, which is not often reported in great detail in motor BMI studies, despite its obvious and important relationship to neural modulation in motor cortical areas.

CHAPTER TWO: Neural and behavioral correlates of learning a beta band neurofeedback task

Abstract

Motor brain machine interface (BMI) studies often assume the subject's intention to move during rigidly structured trials and do not consider "rest" periods. Medical translation of motor BMIs to patient populations relies on successful asynchronous control, in which the patient is able to exert choice about when to control a BMI and when not to. We investigated the possibility of asynchronous control through a beta sensorimotor (SMR) rhythm brain "switch" in a neuromodulation task.

Three mice successfully modulated vibrissal motor cortex (vM1) beta band power to control a 1D auditory cursor. Task performance was due to neuromodulation success and not to advantageous reward seeking responses. Reward threshold rates increased and latencies decreased, occurring soon after trial initiation. Task-related beta SMR activity consisted of bursting events, in agreement with several recent studies. Paralleling the neuromodulation learning was an increase in beta burst durations from 157 ms to 182 ms, suggesting they contributed to performance. Increased whisker pad electromyograms (EMG), and consistent multiunit activity and local field potential phase relationships between vM1 and contralateral EMG during task-related periods, indicate beta power modulation occurred during active motor control. Though the task did not explicitly

require specific movements, the consistency of the vM1-EMG relationships across all three mice suggest the increased burst durations may therefore be a direct result of increased motor activity durations. Further, the vM1-EMG phase coherence increased around reward threshold, at frequencies in the high beta range that are higher than normal mouse whisking frequencies are reported to be, suggesting a tonic "hold" beta was employed by all three mice to solve the neuromodulation task.

The similarity with which mice solved the neuromodulation task suggests that only a subset of beta rhythm phenomenology can be volitionally modulated (e.g. the tonic "hold" beta), in this case through intentional extension of motor processes resembling a "hold". If true, the possible set of successful beta neuromodulation strategies would therefore be limited. While we showed the feasibility of a brain "switch" for asynchronous BMI control, future neuromodulation studies should consider the underlying mechanisms of control, even within a particular rhythm.

Introduction

Brain-machine interfaces (**BMIs**) have shown remarkable progress in the last 15 years (Taylor et al., 2002; Carmena et al., 2003; Guenther et al., 2009; Gilja et al., 2012), and continued progress could one day help paralyzed patients regain lost motor function (Collinger et al., 2013). However, almost all BMI studies use a paradigm of rigidly structured trials that disregard the motor state of the patient, assuming that a patient intends to move only if a trial is ongoing ("synchronous BMI"). While structured experiments have proven helpful to the progress BMI technology has made, translating the resulting decoders into real world applications would have limited benefits if the subject cannot exert choice over when to control the BMI and when not to ("rest periods"; e.g. when sleeping, reading, or conversing). Many BMI decoders show significant drift during rest periods (Chao et al., 2010), as the processing of decoding neural activity never ceases and BMIs interpret slight movement even when there is no intention to move. The inability of BMI decoders to distinguish between motor preparation and execution can further interrupt natural motor dynamics and force accommodations from the BMI subject (Canolty et al., 2012). Before motor BMIs can be translated from the laboratory, it is critical that they be able to determine the neural state of a subject to avoid improperly decoding non-existent motor commands (Ryu and Shenoy, 2009). Only when BMIs are usable outside of a rigid trial structure ("asynchronous BMI"), will their full potential as medical devices be realized (Hochberg et al., 2012).

Unit activity has been used to detect changes in motor state (Achtman et al., 2007; Kaufman et al., 2014; Velliste et al., 2014), but long term unit activity is less stable than local field potentials (Chao et al., 2010; Chestek et al., 2011; Prasad et al., 2012; Flint et al., 2013). Further, LFPs contain more information than unit activity about behavioral state when few electrodes are available for recording (Bansal et al., 2012; Flint et al., 2012), potentially due to a more consistent relationship between LFPs and states (Pesaran et al., 2002; Scherberger et al., 2005; Hwang and Andersen, 2009). LFP-based motor state detection has previously been demonstrated (Mason and Birch, 2000; Bai et al., 2008; Fatourechi et al., 2008; Wang et al., 2012; Williams et al., 2013), and a potential LFP feature that has successfully been modulated in BMIs is the beta rhythm (Wolpaw and McFarland, 2004).

The beta rhythm is a classic neural signature of sensorimotor activity (Sanes and Donoghue, 1993; Murthy and Fetz, 1996) that appears prominently in the entire cortical-basal ganglia loop (Cassidy et al., 2002; Kühn et al., 2004) and could be a natural indicator of motor intent. Event-related synchronization (**ERS**; an increase in power) of beta occurs as rebounds after movement (Pfurtscheller and Lopes da Silva, 1999; Cassim et al., 2001; Pfurtscheller et al., 2005), and occurs during real and imagined movement preparation and planning (Shenoy et al., 2003; Scherberger et al., 2005; O'Leary and Hatsopoulos, 2006; Leeb et al., 2007; Kilavik et al., 2012; Torrecillos et al., 2015). These ERS events seen in trial-averaged data may, however, be constituted by an increased rate of bursting events, rather than sustained oscillatory events (Feingold et al., 2015; Jones, 2016; Lundqvist et al., 2016; Sherman et al., 2016). Therefore, volitional modulation of

the beta sensorimotor rhythm (**SMR**) or these bursting events could be used as a "brain switch" by the subject to indicate when the BMI decoder should start or stop decoding (Pfurtscheller and Solis-Escalante, 2009).

A potential problem with using the beta SMR is that there may be a reduction in SMR activity for paralyzed and "locked-in" patients (Kübler and Birbaumer, 2008). It may also be important to distinguish between motor imagery and motor execution (Birch et al., 2002), which can be difficult because LFP dynamics underlying imagined movement parallel the dynamics underlying actual motor movement (do Nascimento et al., 2006; DaSalla et al., 2009; Miller et al., 2010). Neurofeedback training could help increase, restore, or maintain reduced beta SMR activity or assist in the differentiation of real versus imagined motor intent (Toppi et al., 2014). Further, if beta band activity is vital to motor learning (Feingold et al., 2015; Torrecillos et al., 2015; Cao and Hu, 2016; Tan et al., 2016), then successful neuromodulation of this activity could be important for patients recovering from stroke or other trauma (Daly and Wolpaw, 2008).

Neuromodulation has been performed by rats, cats, mice, monkeys, and humans in cortical and subcortical areas (Wyrwicka and Sterman, 1968; Wolpaw and McFarland, 2004; Cerf et al., 2010; Kobayashi et al., 2010; Philippens and Vanwersch, 2010; Ludwig et al., 2011; Koralek et al., 2012; Sakurai and Takahashi, 2013; Grosse-Wentrup and Schölkopf, 2014), but little of this work details the behavior and neural-behavioral relationships accompanying the neuromodulation.

We investigated the use of beta oscillations in motor cortex as a "switch" for neural control of a BMI, due to the beta rhythm's natural association with motor control, as well as the extended lifespan and usefulness of LFP signals over spiking activity (Wang et al., 2014). Freely moving mice in a neurofeedback task received water reward by controlling an auditory pitch cursor. The cursor was controlled through a modified beta power measure from intracortical electrodes in vibrissal motor cortex (**vM1**). Beta power (13-30 Hz) was normalized by broadband activity to reduce motion artifacts (Ludwig et al., 2011). We were able to investigate the features of beta neuromodulation learning because the electrodes were intracortical. Recently, beta band activity has been found to occur in bursts, rather than sustained oscillations, in both the cortical and subcortical sensorimotor system (Feingold et al., 2015; Sherman et al., 2016) and in prefrontal areas involved with working memory and attentional control (Sacchet et al., 2015; Lundqvist et al., 2016). If beta band activity occurs in bursts, it is unclear how learned neuromodulation of beta activity would alter the power, duration, or rate of beta bursts.

Neuromodulation performance improved quickly, increasing suddenly and plateauing in the 3rd control session. These performance improvements were due to neuromodulation success—not through learned behavioral strategies. Post-hoc analysis of LFP data shows that the neural activity underlying the performance improvement were ~5 – 27 Hz beta-range bursts that increased in median duration from 157 – 182 ms during task-related periods across the experiment. Beta-range bursts did not increase in power or rate during these same task-related periods. Increases in EMG power from ~12 – 38 Hz around reward threshold periods were associated with increases in task-related vM1 LFP-EMG

coherence and significant MUA-EMG phase preferences. The increases in EMG power and apparent drive of high beta EMG by vM1 suggest the possibility that the beta-range bursts resemble the tonic “hold” beta SMR previously described in monkeys and humans. This work illustrates the feasibility of using neurofeedback to create a beta-range “brain switch” for motor BMIs, but also provides evidence that neuromodulation studies should be mindful of the context in which rhythms are modulated.

Methods

Overview

All procedures were performed under protocols approved by the Boston University Institutional Animal Care and Use Committee guidelines. Freely moving mice received instantaneous neurofeedback of left vibrissal motor cortex (**vM1**) beta power and increased beta power relative to broadband power for water reward.

Subjects

Mice in this experiment came from homozygous and heterozygous Emx-Cre (Jin et al., 2000) and Ai32 (Madisen et al., 2012) lines, or subsequent cross-breeding (n = 5 males). While the genetic modifications in these lines did not play a direct role in the experiments, these lines were used in anticipation of future experiments employing optogenetic stimulation. One of the five mice had a high rate of false alarms during the behavioral training phase and was never implanted for neurofeedback control. A second mouse had a relatively long recovery time after the implant surgery, and did not initiate many trials during neurofeedback sessions. We therefore curtailed this mouse's participation in the experiment and its data is excluded from analysis. Data from three mice therefore remained for analysis. These mice will hereafter be referred to as mice A (Emx-Cre +/+), B (Emx-Cre –/+), and C (Ai32 –/+ Emx-Cre –/+), and were respectively P140, P127, and P85 at the start of training.

Mice were maintained on a 12 hr light/dark schedule and given *ad libitum* access to food. Water was restricted the day before experimental sessions and continued

throughout a contiguous session block. Mice received up to 1 mL of water during the course of an experimental session through water rewards, and 1 mL of water several hours after the termination of each experimental session.

Behavioral Task

The arena (Figure 1A) was an open field with two adjacent nose ports: (a) a reward port for water delivery, and (b) a trial initiation port, that mice used to initiate trials (Erlich et al., 2011; Leventhal et al., 2012). In preliminary work, we found that mouse-initiated trials stabilized motivation to perform and increased task participation. Further, mouse-initiated trials allow us the possibility of recognizing mouse intention to participate on a single trial basis.

Upon trial initiation, the house light turned off and auditory feedback simultaneously started, indicating an ongoing trial (Figure 1B). If the beta power reward threshold was not met within 30 s (Koralek et al., 2012) or the mouse committed a false alarm and poked the reward port early, the trial ended and the house light turned on and feedback simultaneously stopped. If the reward threshold was met within 30 s (success), feedback stopped and a 10 kHz auditory tone (reward tone) played for 2-5 s. If the mouse poked the reward port while the reward tone was playing, water reward was provided and the house light turned on. If the reward tone stopped before the mouse poked the reward port, the house light turned on and no reward was provided.

Auditory feedback was presented, in which the frequency of the tone was graded by a logistic function of the beta power ratio magnitude (see *Decoding Features*). EMG from the whisker pads was measured bilaterally.

We define the behavioral index (**BI**) as:

$$I_{beh} = \frac{TP - FA}{TP + FA}$$

where TP is the number of true positives (i.e. rewards received), and FA is the number of false alarms. The range of this value is [-1, 1], with -1 indicating all false alarms, and 1 indicating all true positives.

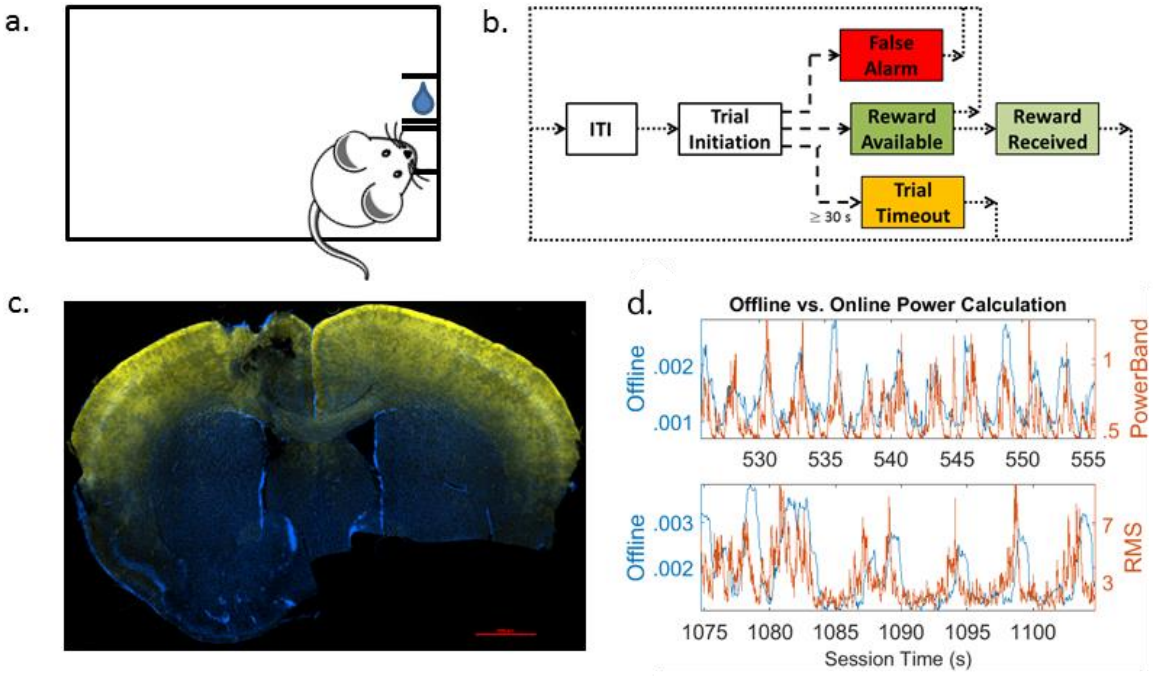

Figure 1: Description of Experiment

(a) Behavioral arena schematic. Two nose ports control the flow of the behavioral task: the trial initiation port (bottom port) and the reward port (top port). (b) Task flow. Inter-trial interval (ITI) periods ended when mice initiated trials by breaking an infrared beam in the initiation port. Upon trial initiation, the neurofeedback auditory tone was played. If mice entered the reward port prior to reward availability, the trial resulted in a

false alarm (red box), the trial was ended, and the ITI period began. If mice successfully reached reward threshold (dark green box), a 10 kHz reward tone played and they had 2-5 s to retrieve the water reward (light green box) before the ITI period began. If 30 s elapsed without a false alarm and without reward availability (yellow box), the trial ended, feedback stopped, and the ITI period began. (c) Histology from mouse C shows electrode tracts in layer 5 of vM1. This slice shows the location of the reference electrode and is located approximately 0.86 mm anterior to bregma. Red scale bar represents 1 mm. (d) Offline BPR calculation confirms accuracy of online BPR calculation. Top: an example of the corresponding offline (blue) and online (orange; PowerBand, i.e. exponential filter) BPR calculations. Bottom: an example of the corresponding offline (blue) and online (orange; beta bandpass RMS) BPR calculations. Note the ranges for both the PowerBand and RMS online calculations do not match the offline BPR calculations, but that the heights of the peaks are similar in scale.

Microelectrode Arrays and Implantation

After mice were trained to seek water reward in response to a 10 kHz auditory tone after ignoring a random auditory tone (1-5 kHz), they were implanted in left vibrissal motor cortex with a microelectrode array (**vM1**; -1.0 mm ML of bregma, +1.0 mm AP of bregma) at an approximate depth of 800 μm to target layer 5 (Paxinos and

Franklin, 2001; Parikh et al., 2009; Matyas et al., 2010; Xu et al., 2012). The arrays were custom 13 channel fixed electrode arrays (approximately 1.5 mm diameter array, 50 µm Teflon coated tungsten wires, ~100-150 kΩ impedance). The reference wire was ~1 mm longer than the other 12 wires, targeted to the subcortical white matter. Three ground wires were wrapped around steel screws (00-90 x 1/16, Antrin Miniature Specialties, Inc.) attached to the skull over the cerebellum, and the left and right cortices. The arrays were then fixed to the ground screws using dental cement.

EMG stereotrodes were created from twisted strands of two 50 µm (100 µm Teflon coated) stainless steel wire (~1 MΩ impedance). One of the wires was cut 1 mm shorter to create a potential differential. The stereotrodes were threaded subcutaneously into the bilateral whisker pads, as described previously (Schroeder and Ritt, 2013).

Histology

Animals were sacrificed and perfused with 0.1 M phosphate buffer and then 4% paraformaldehyde solution. Brains were extracted, placed into a 30% glucose solution, and then sliced at 60 µm. The slices were treated with 1% Triton X-100 detergent, stained with DAPI, and fixed on slides with ProLong Gold antifade reagent. Images were captured using a Nikon Eclipse Ni-E fluorescent microscope and processed with Nikon NIS-Elements Advanced Research software.

Histological slices confirm the placement of the recording electrodes within deeper layers of motor cortex, especially vM1, and the reference electrode below motor cortex within the white matter or lateral ventricles (Figure 1C). Some cortical damage was visible, but due to the length of time (~3 – 6 months) between implantation and

sacrifice, it is unclear to what extent the cortex was damaged during the time of data collection for this work. Based on the presence of MUA (see *MUA Activity*), we believe the cortex to have been reasonably healthy during the relevant recording period.

Neural Recordings

All neural data were recorded using a TDT RZ2 BioAmp Processor and a TDT PZ2 Preamplifier (Tucker-Davis Technologies) at a 24.4 kHz sampling rate. All non-EMG data were low pass filtered at 150 Hz using a 4th order non-causal Butterworth filter and downsampled to 488 Hz for offline processing of LFPs, unless otherwise noted. EMG was processed as follows: (1) the two signals from the stereotrode were subtracted from each other, (2) this differential signal was band pass filtered from 200-3000 Hz using a 2nd order non-causal Butterworth filter, (3) this band passed signal was low pass filtered at 40 Hz using a 2nd order non-causal Butterworth filter, and (4) the resulting signal was downsampled to 244 Hz for offline processing.

At the start of each recording session, "bad" electrodes were identified via visual and auditory inspection and were excluded from online neurofeedback control. All remaining electrodes were considered "good". Only mouse C had any "bad" electrodes for the brain control sessions analyzed (n=1 – 4 electrodes per session). All 12 of mice A's and B's recording electrodes were considered "good" for these analyses.

Recording Sessions

Two different online power calculation methods were used to determine real time BPR values (see *Decoding Features*). Mice A and B began using the first method

(exponential filter) for 12 brain control sessions before switching to the second method (RMS filter). Mouse C used the second method exclusively. The important comparison to be made in this work is the development of neuromodulation skill by novice subjects, and not how subjects use particular power filters. Therefore, to help ensure a fair comparison of neuromodulation skill development between the three mice, we only consider those first 12 brain control sessions for mice A and B in which the exponential filter was used.

Rodents show relatively little desire for water within 24 hours of water deprivation (Stellar and Hill, 1952; Dufort and Abrahamson, 1966). Even when over-trained in the neuromodulation task, the three mice initiated significantly fewer trials during the first session in a contiguous block of session days (median 169 vs. 222, $p = 0.0014$, Wilcoxon rank sum test, Holm-Bonferroni corrected) and significantly more timed out trials per trial initiation occurred during these first sessions (median 0.09 vs. 0.02, $p < 0.001$, Wilcoxon rank sum test, Holm-Bonferroni corrected). A lack of motivation during recent water deprivation is therefore a potential confound and we excluded these initial sessions from further analysis (2 sessions each for mice A and B). After exclusion, there were only ten sessions remaining for the first method (exponential filter) for each of mice A and B. All analyses are based on these first ten, non-initial sessions for the first two mice, and the first ten sessions of the third mouse using the second method (RMS filter).

Decoding Features

Previous work has used a simple beta power threshold in a sliding window to signal motor intent (Leeb et al., 2007). However, the subject in Leeb et al. was a

stationary human who was confined to the wheelchair, whereas our task involves freely moving mice. As discussed in previous studies (Ludwig et al., 2011), motion artifacts are possibly confounds and we therefore corrected for broadband power increases by using a beta power ratio (**BPR**). BPR is defined as the power in the 13-30 Hz band (i.e. beta) normalized by the power in the 1-80 Hz band (broadband LFP), similar to previous normalization strategies (Hamada et al., 1999; Shaw and Chew, 2003; Kropotov et al., 2005; Sun and Dan, 2009; Rouse et al., 2013; Williams et al., 2013; Khanna and Carmena, 2015).

We used two methods of calculating the power in the beta band. The first method used TDT's PowerBand module, which calculates instantaneous power using exponential filters of the cosine and sine of the neural data. The second method was bandpass filtering the raw neural signals, and then calculating the RMS of these filtered data. We will refer to the PowerBand module method as the "exponential filter", and the RMS method as the "RMS filter".

The PowerBand module calculates power via the following equation:

$$S_t = cosAvg_t{}^2 + sinAvg_t{}^2 + 2|(cosAvg_t * sinAvg_t)|$$

where $cosAvg_t$ and $sinAvg_t$ are exponential moving averages of the cosine- and sine-multiplied values of the neural data:

$$cosAvg_{t+1} = (1-\tau) * cos_t + \tau * cosAvg_t$$

$$sinAvg_{t+1} = (1-\tau) * sin_t + \tau * sinAvg_t$$

where τ is the exponential weighting of the average, determined by the bandwidth B and the sampling frequency Fs:

$$\tau = e^{-2\pi * B * Fs}$$

and cos_t and sin_t are the instantaneous products of the neural signal and the cosine and sine functions:

$$cos_t = x_t \cos(2\pi\varphi t)$$

$$sin_t = x_t \sin(2\pi\varphi t)$$

where φ is the center frequency of the band.

For mice A and B, we used three PowerBand modules to calculate power in the 0.5-55.5, 64-80, and 14-30 Hz bands ($\varphi = 25.5$, $B = 50$; $\varphi = 72$, $B = 16$; and $\varphi = 22$, $B = 16$, respectively). The 14-30 Hz band power was divided by the sum of the power in the 0.5-55.5 and 64-80 Hz bands. For mouse C, we used the RMS of the 1-57, 63-80, and 13-30 Hz causally filtered neural data (12 dB/octave filter rolloff).

A single average BPR was calculated from the BPR of all individual electrodes. This average BPR was smoothed by 600 ms (Koralek et al., 2012) for the PowerBand calculation, and the RMS values were smoothed by 1000 ms. We also compared both online BPR calculation methods (exponential and RMS) and an offline version that calculated power through the multitaper method (chronux.org) and was then summed and normalized similarly to the online BPR estimate (Figure 1D). The online BPR calculation is noisy and has a different range than the offline calculation, but has a similar relative scale and peaks at similar times.

BPR Threshold Determination

After implantation, mice were given approximately a week to recover from surgery and were then recorded during pre-brain control sessions to determine the

appropriate value of the BPR reward threshold. Mice A and B each ran two pre-brain control sessions. Due to recording quality issues, we ignored the first pre-brain control session for both mice A and B. BPR threshold values were determined from the remaining pre-brain control session for each mouse. Mouse C ran five pre-brain control sessions and we determined a BPR reward threshold value as the average BPR threshold of these sessions.

We individually calibrated the BPR threshold offline for each mouse using the pre-brain control session BPR data. We iteratively adjusted the BPR threshold value until mice would have reached reward threshold 80 times in a 40 min session (i.e. two rewards/min). This initial rate of reward was found to be appropriate in initial modeling because it was low enough to allow improvement and not so infrequent to demotivate mice unfamiliar with the brain control task.

Neurofeedback

Auditory tone frequency was altered in real time, according to:

$$Hz_t = \frac{4040}{1 + e^{-\left(4.6(BPR_t - BPR_{mean}) / (BPR_{thresh} - BPR_{mean})\right)}} + 1000$$

where Hz_t is the auditory frequency at time t, BPR_t is the online BPR at time t, BPR_{thresh} is the BPR reward threshold for the mouse, and BPR_{mean} is the mean of BPR_{thresh} and the 1st percentile of the BPR values found from the pre-brain control sessions. This logistic function bounds the feedback frequency between 1000 Hz as $BPR \to \infty^-$ and 5040 Hz as $BPR \to \infty^+$, and 5000 Hz (reward target) at $BPR_t =$

BPR_{thresh}. In preliminary work, we found a logistic function necessary because the limited value range allows the same parameter values to be used over the entire experimental life of a mouse, regardless of any signal quality changes or BPR drift.

MUA Activity

Multi-unit activity was extracted through a two-step process. First we applied a bandpass filter to the raw neural recordings between 300-5000 Hz using a non-causal 6th order Butterworth filter. The bandpassed neural activity was then inverted so that only negative-going spikes were detected. MUA windows (~1.5 ms) were selected using the following threshold (Quiroga et al., 2004):

$$Thr = 4\sigma_n$$

$$\sigma_n = median\left\{\frac{|x|}{0.6745}\right\}$$

where σ_n is an approximate standard deviation that reduces the effect of high spike amplitudes and bursting or fast-firing units on the standard deviation, and x is the bandpassed neural activity. MUA windows were then aligned to the maximum value of the window, in contrast to aligning on the threshold crossing times.

Second, we performed the following three quality control steps on the extracted MUA windows to exclude MUA activity that was likely due to noise:

MUA activity that occurred during LFP signal saturation ($|LFP_x| > 0.8\ LFP_{MAX}$) was considered to be noise and was excluded.

MUA activity with a full-width half-max (**FWHM**) ≥ 10 samples (~0.41 ms) had shapes that appeared to be noise and were excluded.

MUA activity was binned in 1 ms bins across all 12 recording electrodes, and any MUA windows that occurred in the same 1 ms bin across ≥ 6 electrodes were considered to be noise and were excluded. To determine that this step was necessary, we permuted the ISIs of all MUA windows across all 12 electrodes 100 times to generate 95% confidence bands of random spiking activity and determined that coincident spikes should not generally occur in greater than ~2 electrodes, again assuming the spikes are completely random. Given that spiking activity is not completely random, we made the cutoff for coincident MUA windows 6 electrodes, to allow for the possibility that some coincident spikes are due to synchronization of neural activity or an increased firing rate. The mean MUA shapes appear to be neural and the retention rate of MUA windows after these quality control steps were 73.3%, 78.6%, and 54.9% of all detected MUA windows, respectively for the three mice.

Burst Analysis

We used methods similar to previous studies to determine the existence and duration of beta band bursting activity (Feingold et al., 2015; Lundqvist et al., 2016). First, we chose frequencies appropriate for each mouse, based on the largest power increases seen 0 – 250 ms prior to reward availability (16-25 Hz, 10-27 Hz, and 5-17 Hz for mice A, B, and C, respectively). Then we band passed the raw neural data using a 4^{th} order non-causal Butterworth filter and these mouse-specific frequencies. Next, we calculated the Hilbert envelope of the band pass filtered data, and defined the existence of bursts where the envelope exceeded 3 times the median value of the envelope, and the duration of those bursts as the continuous data above 1.5 times the median value.

Results

Mice Learn the Neuromodulation Task

Mice successfully learned the neuromodulation task and several lines of evidence suggest mice accomplished this feat through proficiency in neuromodulation and not to advantageous behavioral strategies such as delayed reward seeking responses.

Raw LFP traces (Figure 2) indicate that there are two major variations of activity that occurred at reward threshold when the online decoder calculated a high BPR: (1) a visually classical beta oscillation, in which the hallmark arch shape repeats at 13-30 Hz, and (2) large negative impulse-like activity, which could be a local motor potential (Schalk et al., 2007). Spectrograms of the neural data around reward availability periods show a power increase in frequencies overlapping the 13-30 Hz range, compared to random baseline periods (Figure 3A). We also saw a progressive increase in the 13-30 Hz range in session-wide LFP spectra (Figure 3B) that could be due to increased numbers of beta band bursts over sessions.

Mice successfully learned the neuromodulation task, as measured by the behavior index (**BI**). The rising BI indicates a decrease in the number of false alarms committed relative to the number of true positives over the ten brain control sessions (Figure 4A, top). The mean BI for all mice begins at -0.5 (i.e. 25% true positives) and peaks at $+0.35$ (i.e. 68% true positives) during session 9. There was a large increase in mean BI at session 3 that could be due to reward seeking strategies and/or successful neuromodulation—inhibition of reward seeking at inappropriate times would inherently decrease the number of FAs, and successful neuromodulation would increase the chances

that reward seeking occurs during reward availability. However, neuromodulation performance is likely the proximal cause for the BI jump at session 3, based on several reward threshold timing measurements that improved during the same session.

The most direct evidence for the neural origins of task success is that the rate of threshold events quickly increases and their timing is tied to trial initiation. Mice quickly increased the rate of BPR threshold events (Figure 4A, middle), irrespective of whether the threshold event occurred within a trial (i.e. whether or not a reward was associated with the event). With only one exception (mouse B, session 6), all mice had higher threshold event rates from session 3 onwards than in either of the first two sessions, which is a direct result of successful neuromodulation and tightly corresponds with the BI jump at session 3. Mice may have learned to increase overall levels of motor activity to generate more reward threshold events at random times. However, we found that threshold events occurred shortly after trial initiation (Figure 4A, bottom). We determined that the post-initiation threshold event latency distributions were not random by shuffling inter-trial interval times and finding the successive reward threshold events. By breaking this time dependence on the actual trial initiation times, we found that all three mice had a significant increase in reward threshold latencies sooner than 2 s after trial initiation, and that the shuffled reward threshold latencies looked more like an exponential distribution that is indicative of randomly occurring events (Figure 4B). Therefore, it seems probable that neural strategies are at least partially responsible for the increase in neuromodulation task performance, but did the mice engage in any reward seeking strategies that could also have increased performance?

Mice could have adopted a strategy of simply waiting long enough for a threshold event to occur before seeking reward, that means we would expect the false alarm (**FA**) latency to be significantly shorter than the true positive (**TP**) latency. In other words, false alarms would be caused simply because mice sought reward too quickly, and so the true positive latency would be longer, on average. If the TP and FA latencies are similar, however, then the reward seeking responses are relatively random and any increase in the number of TPs relative to FAs must be due to the ability of mice to generate threshold events soon after trial initiation. We found the latter relationship to be true: median latencies to TPs and FAs from trial initiation were similar across most sessions for all mice (Figure 4C). This suggests that mice tended to seek reward at random intervals and that the success rate depended on the latency to reward threshold from trial initiation. Moreover, both the TP and FA latencies decreased over the first 3-4 sessions, and then generally remained constant for the remaining sessions, arguing against the possibility that TPs increased and FAs decreased because the mice simply waited longer to respond. TP and FA responses were similar enough, however, that we wondered whether mice were seeking reward entirely randomly or whether it at least partially depended on auditory feedback frequency. If TP response latencies were completely random, the probability density should resemble an exponential distribution. However, mice reliably sought reward ~250 to 650 ms after reward tone, with a notable decrease in response probability ~0 to 250 ms after reward tone. These data indicate that reward seeking was a response to the reward tone.

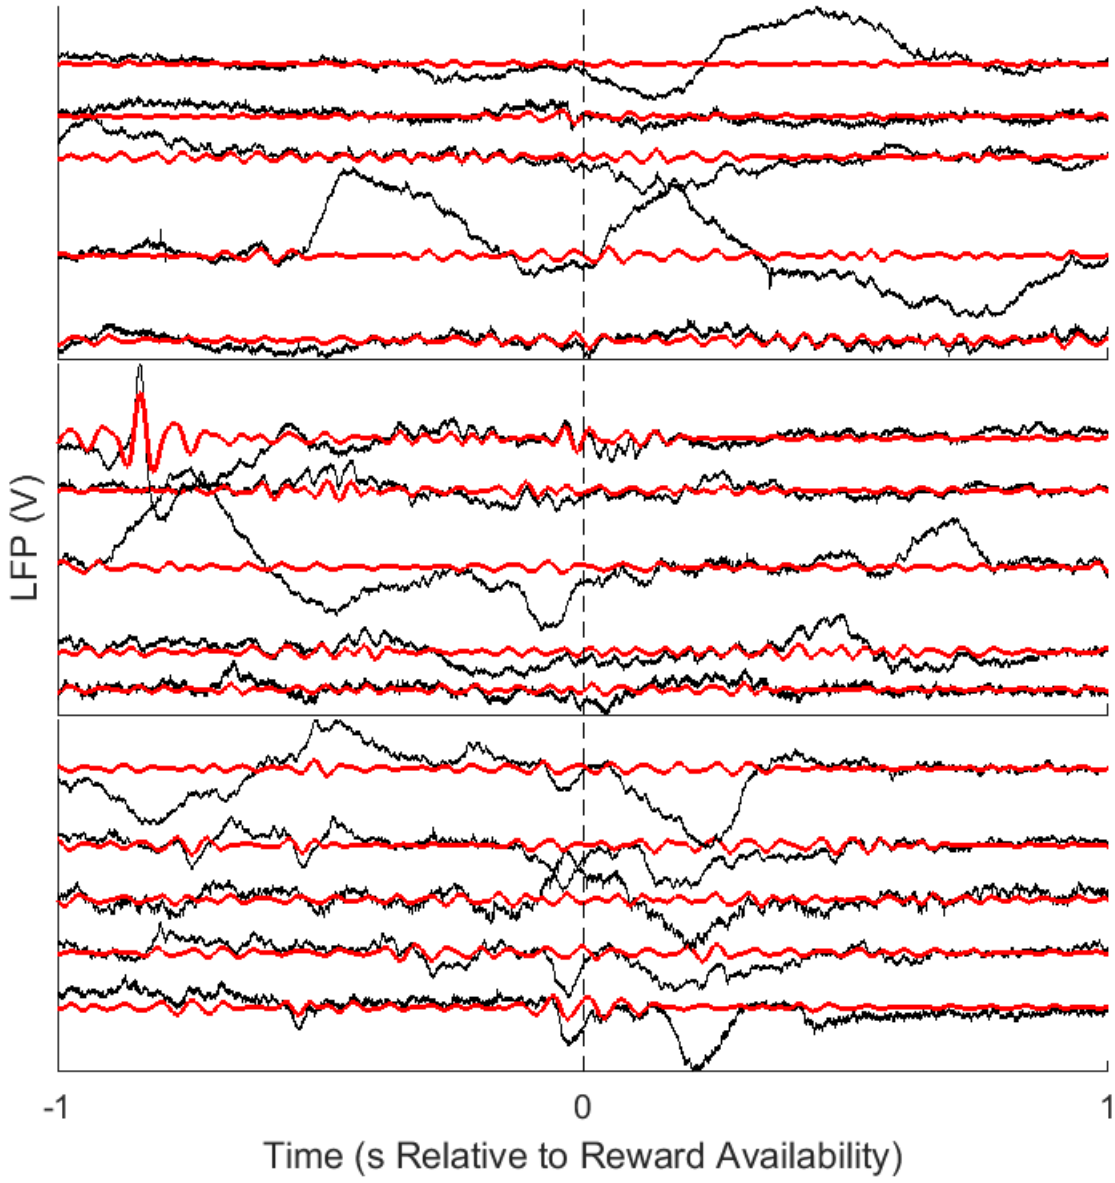

Figure 2: Raw LFP examples

Five random LFP traces (black) are shown for each of the three mice (top to bottom, respectively). Each trace is centered on the 2 s around reward availability. The bandpassed LFP traces (13-30 Hz; shown in red) are plotted for comparison. There are varying degrees to which classically-defined beta band activity occurred. For example, most of the raw traces

for the first two mice (A and B) have many bouts of classical beta band activity, including true 13-30 Hz oscillations and the hallmark arch shape. In contrast, the third mouse (C; bottom) had little classically-defined beta band activity, but instead had only brief moments in a true 13-30 Hz regime, without the appropriate arch shape. Also note that there may not be large 13-30 Hz oscillations on a given electrode around reward availability, but that there would be ~11 other electrodes with concurrent activity that may have had larger 13-30 Hz oscillations that increased the BPR value at that time.

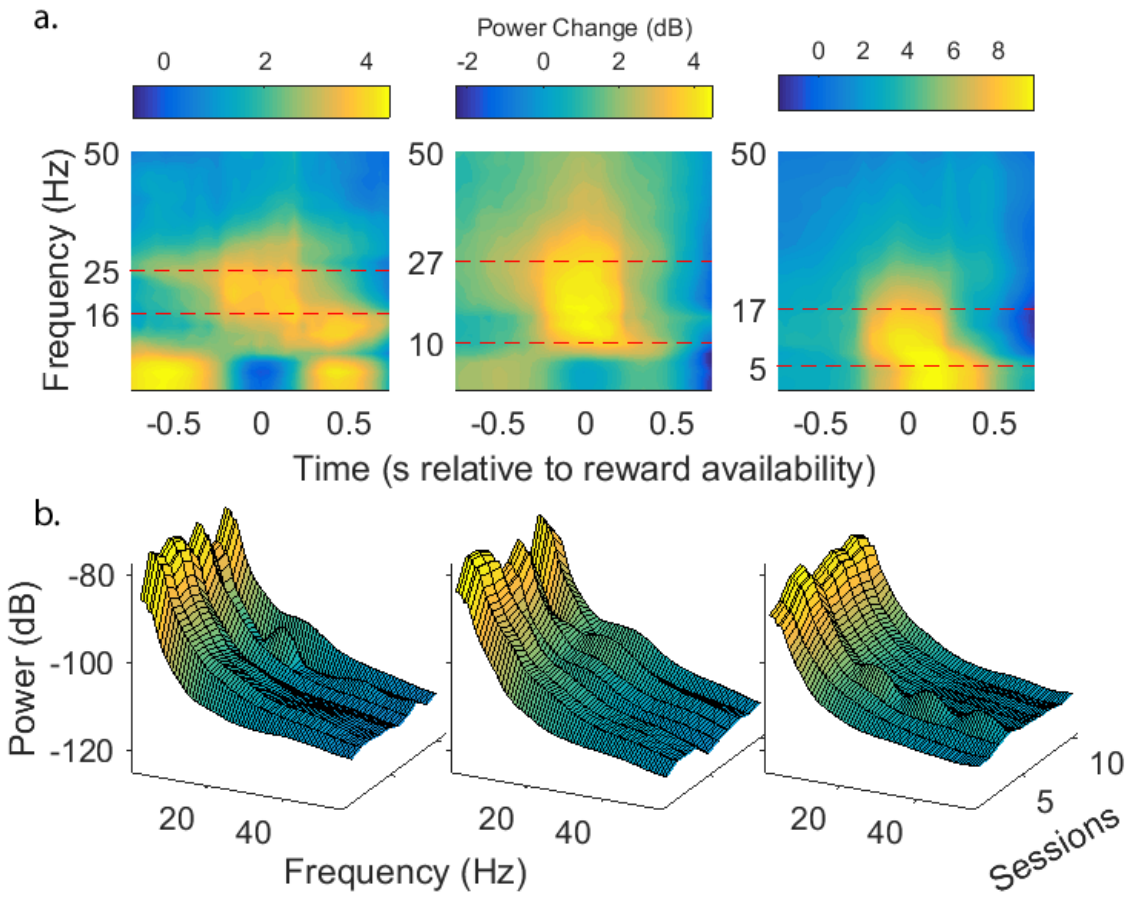

Figure 3: Beta band power increased over sessions

(a) Mice showed an overall increase in ~5 – 45 Hz power around reward availability. All spectrograms are normalized by random epochs. Red dashed lines indicate frequency bands in which there was a large increase in power 0 – 250 ms prior to reward availability. These subject-specific frequency bands are used for burst analyses (see *Bursts Underlie Beta Band Neuromodulation*). (b) LFP spectra show bumps in the 15 – 30 Hz range that increased over sessions. Spectra are from random epochs.

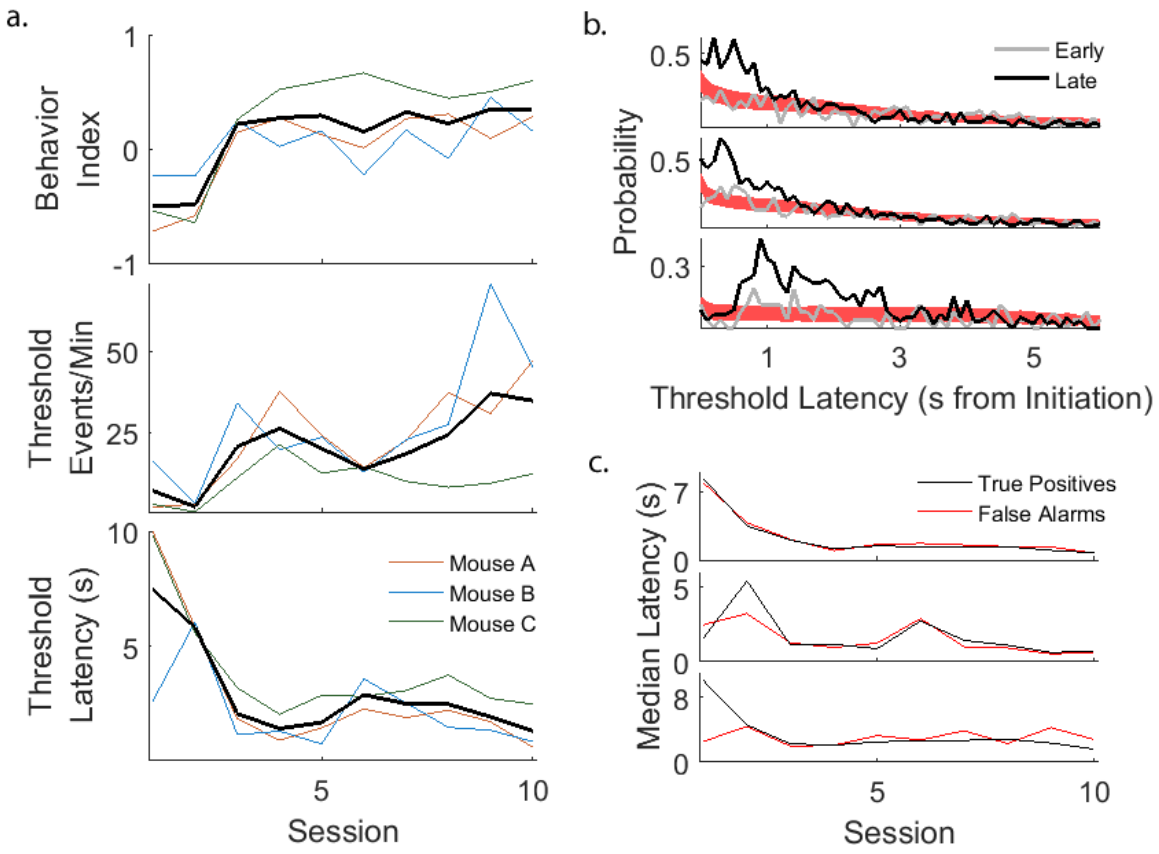

Figure 4: Task performance increases are due to neuromodulation success and not due to reward seeking strategies

(a) Behavior index (BI) increases over sessions (top) are supported by increases in threshold rates (middle) and decreases in reward threshold latencies (bottom). Note the BI jump in session 3 is accompanied by a similar jump in the threshold rate and drop in threshold latency in session 3. This result suggests the BI increase is at least partially due to neuromodulation success. (b) Reward threshold latencies were tied to trial initiation times. Early session (1-3) latencies (gray) occur within a similar distribution as random bootstrapped latencies (red). Late session (8-10) latencies (black) occurred earlier than by chance, suggesting that mice

learned to produce threshold events soon after trial initiations. Bootstrapped latencies were generated by permuting the inter-trial intervals (ITIs). (c) True positive and false alarm latencies decreased over sessions and are not substantially different, indicating that the BI increase was not due to delayed reward seeking responses.

Bursts Underlie Beta Band Neuromodulation

Recent evidence suggests that trial-averaged spectrograms misrepresent the neural activity that occurs at the single-trial basis (Feingold et al., 2015; Lundqvist et al., 2016; Sherman et al., 2016). For instance, trial-averaged spectrograms might suggest prolonged activity in the beta band range during the preparatory period before a reach (Sanes and Donoghue, 1993), when in fact the power increase seen over time periods as long as a few seconds may only be the result of an increase in the rate of short ~100 – 200 ms burst events.

Although the reward criterion of our neuromodulation task was to increase power in the 13-30 Hz band relative to broadband power, our mice generated power increases in various frequencies bands between ~5 – 45 Hz (Figure 3A). We therefore used mouse-specific frequency bands (16 – 25 Hz, 10 – 27 Hz, and 5 – 17 Hz, for mouse A, B, and C, respectively) with increased power over the 250 ms before reward threshold to investigate whether the power increase consisted of bursting events.

To establish that the increased power consisted of bursts, rather than prolonged rhythmic activity, we first looked at the duration of the elevated power events using

methods similar to recent work (Feingold et al., 2015; Lundqvist et al., 2016; Sherman et al., 2016). The median duration of elevated power events across all "good" electrodes in all brain control sessions was 205, 133, and 207 ms for mouse A, B, and C, respectively, which correspond to 3.3 – 5.1, 1.3 – 3.6, and 1.0 – 3.5 cycles per event, depending on frequency. These median durations are similar to previous descriptions of beta bursts (Murthy and Fetz, 1992, 1996; Leventhal et al., 2012; Feingold et al., 2015; Lundqvist et al., 2016) and suggest the possibility that beta range bursts underlie the power increases seen around reward threshold.

We next looked at the density of bursts around both trial initiation and trial termination (when the mouse either commits a FA or TP, or the 30 s trial time limit was reached). There was a large increase in burst density starting ~0.5 s before trial initiation, which dropped to or below chance density for ~0.5 s starting at trial termination, when mice sought reward or had otherwise not yet initiated the next trial (Figure 5). The large increase in burst density from ~0.5 s prior to trial initiation until trial termination suggests bursting activity was enhanced during trial-related time periods. We therefore defined bursts to be "task-related" if their onset occurred between 0.5 s before trial initiation and 0 s from trial termination (Figure 6A). Having defined task-related burst events, we found the median task-related burst durations for two of the three mice were lower than all task- and non-task-related bursts (144, 142, and 184 ms for mouse A, B, and C, respectively, corresponding to 2.3 – 3.6, 1.4 – 3.8, and 0.9 – 3.1 cycles, depending on frequency). Therefore, it is possible that whatever behavioral or neural strategies the mice use to

solve the neuromodulation task may be inherently different than incidental beta burst-producing activities.

We next asked if the neuromodulation performance increase over sessions was due to an increase in the amplitude of task-related bursts, given the explicit reward state of high 13-30 Hz power. The median-normalized RMS of the task-related bursts did not significantly increase over sessions, although the slope approached significance (Figure 6B; $p = 0.09$, permutation test, $p = 0.18$ after Holm-Bonferroni correction). We then wondered if the performance improvement was due to an increased rate of task-related bursts, due to the lack of significant increase in burst RMS and the increase in burst density around trials. Again, the rate of task-related bursts did not significantly increase, although the slope approached significance (Figure 6C; $p = 0.09$, permutation test, $p = 0.18$ after Holm-Bonferroni correction). We also measured the median duration of task-related bursts and found that they significantly increased from 157 ms to 182 ms (Figure 6D; $p = 0.01$, permutation test, $p = 0.03$ after Holm-Bonferroni correction). The 25 ms change corresponds to a median increase of ~0.1 – 0.7 cycles per burst, depending on the frequency (5 – 27 Hz).

These durations agree well with the literature, increase over sessions, and may be a major driver of neuromodulation task success. Both the power and rate of the task-related bursts do not significantly increase, and could indicate that they are not the proximal causes of neuromodulation task success. In sum, the trial-averaged power increases seen around reward threshold appear to be generated by an increase in the

duration and density of task-related burst events, rather than through an increase in the power of extended oscillations.

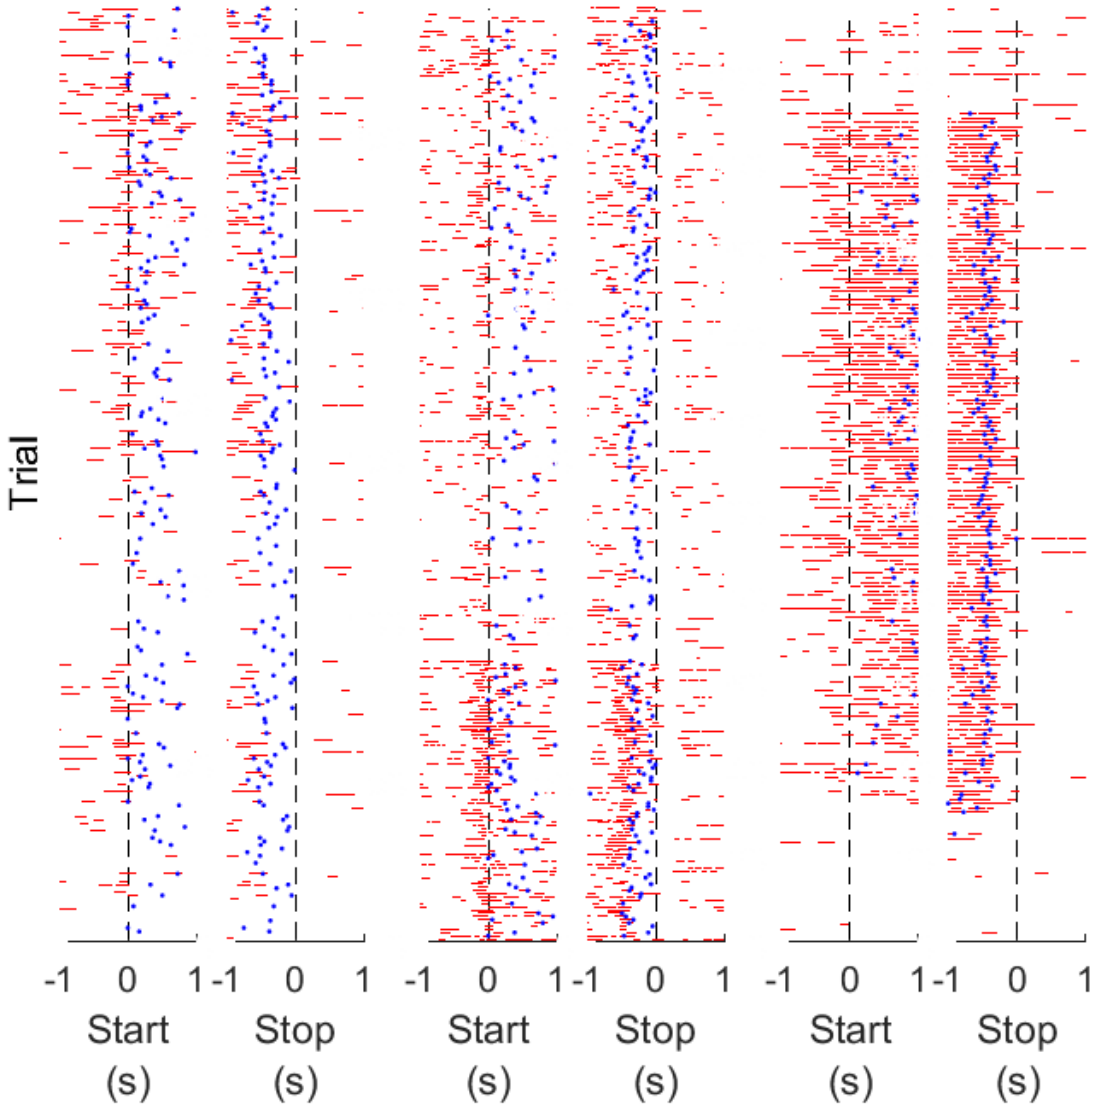

Figure 5: Beta band range bursts underlie beta band power increases

Examples of elevated beta range bursts during the last session of each mouse. Red lines: elevated power event (i.e. a "burst"). Blue dots: reward threshold events. The number of beta bursts increased ~0.5 s prior to trial starts and dropped at trial termination. Reward threshold events occurred

during periods of dense beta bursts. The mice had different progressions of beta burst generation during this session; mouse A gradually increased the production of burst events, mouse B had a fairly sudden drop in the production of burst events, and mouse C had a very distinct period of burst production in the middle of the session. Mice A and B generated reward threshold events throughout their sessions, but mouse C only had reward threshold events during the period of increased burst production. Traces are in trial order (first trial at bottom, last trial at top), and a random "good" electrode was selected for display in each trial. Left plots: mouse A. Center plots: mouse B. Right plots: mouse C.

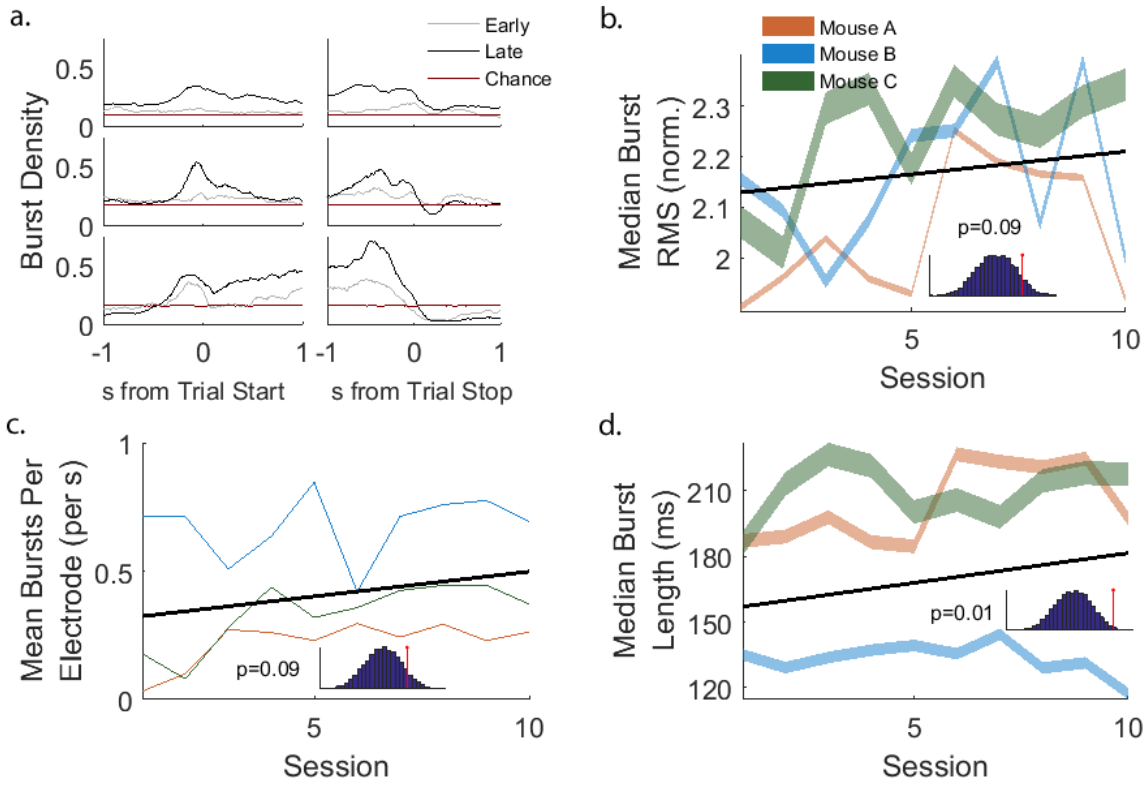

Figure 6: Beta range burst durations increased over sessions

(a) Burst density increased ~0.5 s before trial initiation and decreased around trial termination. This time range defines the period of "task-related" bursts. Early burst densities (gray) were lower than late burst densities (black) and were similar to chance burst densities (red). Chance burst densities were derived from 10000 random epochs from all sessions. Top row: mouse A. Middle row: mouse B. Bottom row: mouse C. (b) The median normalized burst RMS did not increase over sessions (p=0.09 uncorrected p-value, 10000 permutations of Theil Sen fit). Inset: histogram of permuted Theil Sen fit slopes and the slope of the unshuffled data. (c) The rate of burst events did not increase over sessions (p=0.09

uncorrected p-value, 10000 permutations of Theil Sen fit). Inset: histogram of permuted Theil Sen fit slopes and the slope of the unshuffled data. (d) Burst durations increased over sessions (p=0.01 uncorrected p-value, 10000 permutations of Theil Sen fit, p=0.03 Holm-Bonferroni corrected p-value). Inset: histogram of permuted Theil Sen fit slopes and the slope of the unshuffled data.

Beta band neuromodulation is achieved through active motor control

Previous neuromodulation studies noted that animals adopt various motor behaviors as part of their task strategies (Wyrwicka and Sterman, 1968; Gage et al., 2005; Ludwig et al., 2011). Our mice were implanted in left vibrissal motor cortex, leaving open the question of whether they adopted whisking behaviors during the neuromodulation task. Behavioral, video, and EMG data suggest that successful neuromodulation required active task engagement (or an "intention" to participate), and that BPR increases were concurrent with increases in EMG activity and consistent vM1-EMG MUA and LFP relationships suggestive of motor "holding".

BMI performance is dependent on motivation (Musallam et al., 2004). In the present task, mice appeared to be engaged in the task at the beginning of experimental sessions, as seen by the rate of trial initiations (Figure 7A). However, later in those sessions when mice were closer to satiation, the rate of trial initiations decreased and trial time outs began to occur. If TPs were dependent on behavioral response strategies rather than neuromodulation performance, the trial time outs would likely have been TPs instead. Video data also suggest that mouse orientation at the moment of trial initiation indicated whether a timed out trial would occur (Figure 7A). Normalized and averaged video frames taken from the moment of trial initiation in which reward was received show an overall body orientation of the mouse's nose in or near the initiation port and the rear of the mouse aimed away from the initiation port, suggesting that mice intended to initiate trials in which they successfully modulate beta-range activity. In contrast, video frames taken from the moment of trial initiation in which time outs occurred show no

common body orientation on average, suggesting that mice may have accidentally initiated these trials and are not actively participating in the neuromodulation task (anecdotally, timed out trials generally seemed to be initiated with paws, tails, or abdomens). Therefore, successful modulation of beta-range power might have required motivation and some form of intention to participate in the task.

We also analyzed EMG power to determine whether neuromodulation was an active process. Overall EMG activity between the first 3 s of true positive and timed out trials (up to but not including reward availability) was higher for true positive trials (data not shown). Because trial time outs occurred late in sessions when mice were more satiated, this finding suggests ongoing whisker or orofacial movements occurred more often during trials in which mice were more likely to "intend" to participate in the task or to receive reward. We examined EMG around reward threshold and found that there was a temporally-precise and frequency-specific bilateral increase in the 10-38 Hz band starting ~400 ms before reward availability (Figure 7B). The power increase was greater in the right (contralateral) EMG than the left EMG for mice B and C. A relatively flat band in the spectrogram appears between ~8 – 16 Hz over the entire analysis window for all three mice, within normal mouse whisking frequencies (Cao et al., 2012). Foveal whisking occurs at higher frequencies than exploratory whisking (Berg and Kleinfeld, 2003), and a whisker tracking study reported the whisking frequency of freely exploring mice to be ~15 – 25 Hz, with high frequency activity (~37 Hz) at the set points (Voigts et al., 2008). These data could indicate that ~8-16 Hz whisking occurs within broad time periods around reward threshold for all three mice, but that reward threshold conditions

are contingent on neural processes that produce additional EMG activity in the 16-35 Hz range, possibly small amplitude whisking. We therefore suspect that mice may involve whisking and control of beta frequency whisker EMG activity as part of their neuromodulation strategy.

Previous studies noted potential whisk-by-whisk motor control based on vM1-EMG coherence in the rat (Ahrens and Kleinfeld, 2004), and we saw increases in LFP-EMG coherence and MUA-EMG phase preferences that suggest a direct vM1 drive of vibrissa motions. LFP-EMG coherence within ~6 – 16 Hz increased bilaterally in all three mice over the ~1 s before reward availability, but the right whisker pad had much more consistent increases in beta-range (~14 – 38 Hz) LFP-EMG coherence (Figure 7C). Additionally, there was a positive slope of the coherence phase in the beta range for the right EMG (data not shown) that is consistent with vM1 LFPs leading EMG activity and driving coherence, rather than whisker pad afference leading LFPs and driving coherence (Schoffelen et al., 2005). The more consistent increase in right EMG-left vM1 LFP activity and vM1-leading LFPs are consistent with contralateral control of the whisker pads.

We also looked at vM1 MUA and found that many of the 12 electrodes (n=3-10) in left vM1 were significantly phase-selective to bilateral EMG (Figure 7D). MUA with significant EMG phase preferences fell within a consistently narrow range for all mice ($^{-\pi}/_{6} < \theta < 0$). These phases occur just before EMG peak, which corresponds to the onset of whisker protraction during normal whisking behavior (Schroeder and Ritt, 2013). Though the angular dispersion values were low ($r \sim 0.02 - 0.27$), even low spike-

field coherence can indicate significant numbers of units are synchronized with the field (Baker et al., 2003). In our task it may only suggest direct vM1 control of high frequency whisker pad EMG activity—as opposed to whisk-by-whisk control—with a slight contralateral bias due to the increased contralateral LFP coherence.

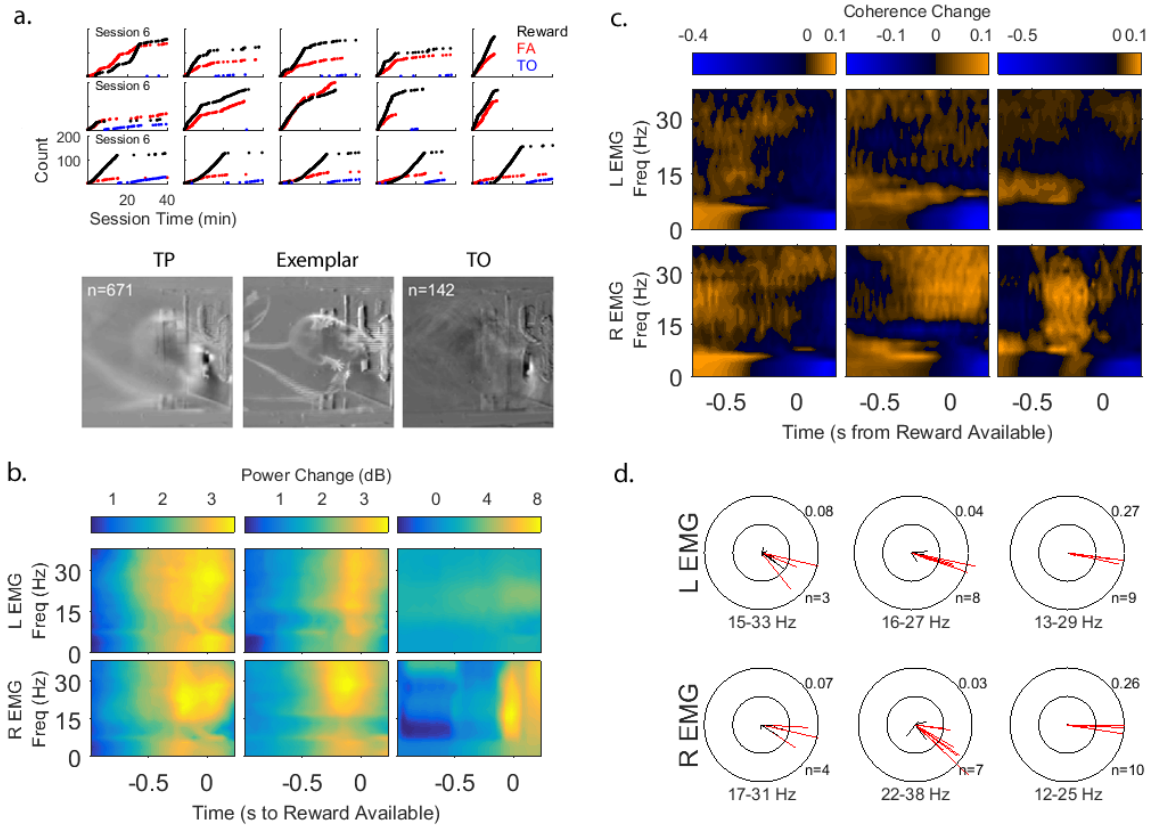

Figure 7: Beta band neuromodulation is an active process driven by vM1

(a) Top: Timed out trials (TO; blue) occur late in sessions, suggesting the involvement of motivation in neuromodulation success. Trials that time out generally occur during plateaus in rewards (black) and false alarms (FAs; red). These plateaus occur after a sharp rise in rewards and/or FAs, indicating the mice slow the pace at which they initiate trials. The reduction in trial initiation and increase in trial time outs suggest that some

level of task motivation or “effort” is vital for beta neuromodulation success. Shown are sessions 6 – 10 for mice A, B, and C (top to bottom rows, respectively). Bottom: Timed out trials occur when mice do not intend to participate in the neuromodulation task. At left is the mean frame at trial initiation for TP trials for mouse B. In the center is the single video frame with the highest correlation to the mean TP trial initiation frame, for comparison. At right is the mean frame at trial initiation for TO trials for mouse B. Note that the general shape and features of the mice are visible in the left frame, and that there is no discernible mouse shape in the frame on the right, despite there being fewer frames in the average (n=142 TOs vs. n=671 TPs). (b) High beta band range EMG power increases around reward threshold. Contralateral EMG power (right EMG; bottom row) increases more than ipsilateral EMG power (left EMG; top row), suggesting a need to increase motor output from left vM1 during beta neuromodulation. The contralateral EMG power increase for mouse C is in lower frequencies than mice A and B, which corresponds well with the lower frequency beta band LFP power increase of mouse C around reward threshold (~5-17 Hz, Figure 3). Left column: mouse A. Center column: mouse B. Right column: mouse C. Power is normalized by random epochs. (c) Beta band range LFP-EMG coherence increases around reward threshold. There is a fairly consistent ~6 – 16 Hz LFP-EMG coherence for all mice, but contralateral EMG coherence (LFP-right EMG; bottom row)

increases more in the high beta band range (~14 – 38 Hz) than ipsilateral EMG coherence (LFP-left EMG; top row). LFP-EMG coherence of random epochs was subtracted from the -1 s to 0.5 s around reward availability. (d) EMG phase preferences of task-related MUA suggest vM1 drives high frequency whisker pad EMG activity. The mean phase preference of all significant MUA phase preferences for both left and right EMG for all three mice lie within a very narrow range ($^{-\pi}/_{6} < \theta < 0$), corresponding to the onset of whisker protraction (Schroeder and Ritt, 2013). Below each plot is the EMG frequency range over which the MUA activity was tested. Red radial lines: statistically significant MUA activity ($p \leq 0.001$, uncorrected Rayleigh z test). Black radial lines: statistically insignificant MUA activity ($p > 0.001$, uncorrected Rayleigh z test). Upper right text (e.g. “0.08”): angular dispersion r value at the outer ring; inner ring is half that value. Lower right text (e.g. “n=3”): number of electrodes (out of 12) for which MUA phase preference was statistically significant.

Discussion

Mice quickly learned the neuromodulation task in the 3rd session, with modest performance gains over the remaining sessions. This timescale of learning is very similar to other studies in which mice modulated neural activity to control an auditory cursor (Koralek et al., 2012; Clancy et al., 2014). Reward threshold LFP power increases occurred mostly in beta frequencies (~5 – 40 Hz), which consisted of an increase in the density of beta range burst events rather than an increase in the magnitudes of extended oscillations. Task performance improvements were likely due to neuromodulation strategies that include increases in duration of the underlying beta range bursts. Contralateral whisker pad EMG activity also increased in the beta range (~12 – 38 Hz) around reward threshold and had high LFP beta coherence and consistent MUA-EMG beta phase preferences.

The coincident increases of LFP and EMG beta power coupled with beta LFP-EMG coherence and MUA beta phase preferences is suggestive of frequency-specific motor patterns consistent with "holding" or other high frequency orofacial motor control. Beta band control in our task and the increased durations of beta bursts may therefore be due to intentional increases in the duration of active motor processes. The rapidity with which mice learned the neuromodulation task and similarity of LFP- and MUA-EMG relationships across all three mice suggests a common, accessible mechanism of beta burst production that motor "holding" or similar cycle-by-cycle orofacial motor control would parsimoniously explain.

Tonic "Hold" Motor Output Driven by vM1

Normally, LFP power in beta frequencies decreases at the onset of movement, rather than increases. One exception to this is the "hold period" beta SMR. Based on LFP, MUA, and EMG evidence, our mice may have learned to actively "hold" their whiskers, similarly to previous studies in monkeys and humans in which beta band power and coherence increases during tonic hold periods and sensorimotor maintenance (Kilner et al., 2004; Omlor et al., 2007; Engel and Fries, 2010). The beta range activity we then see could be static, rather than dynamic, muscle control and the burst event density decrease that occurs just as mice seek reward in response to the reward tone would then correspond to the initiation of dynamic movement when beta power decreases (Donoghue et al., 1998). This interpretation fits with classical descriptions of the beta SMR, as similar beta LFP-EMG relationships have been reported in macaque monkeys (Baker et al., 1997) and humans (Schoffelen et al., 2005; van Ede and Maris, 2013). High gamma coherence (40 – 47 Hz) has also been reported in humans during a bimanual wrist extension task (Schoffelen et al., 2011), which supports the possibility of high frequency vM1-EMG coherence in mice. The consistent MUA-EMG phase relationship we found also supports the possibility of vM1-high frequency EMG control. There is no current consensus as to how vM1 activity relates to vibrissa movements (Castro-Alamancos, 2006; Hill et al., 2011) and EMG activity, and the particular task requirements of this study could have altered normal vM1-whisker pad dynamics. Mice could have generated increases in BPR via motor artifacts (Ganguly and Kleinfeld, 2004) or intrinsic control that is dissociable from movement (Lebedev et al., 2005; Moritz and Fetz, 2011),

although we suspect the latter possibility is unlikely (see below). The high frequency EMG activity (approximately in the beta range) may therefore be related to tonic "hold" activity.

Alternatively, our mice might not have been generating "true" beta SMRs, but were rather controlling whisking or orofacial behaviors on a cycle-by-cycle basis (Berg and Kleinfeld, 2003; Ahrens and Kleinfeld, 2004). This locking could be achieved by a phase reset of the whisking cycle by layer V cells (Brecht et al., 2004). Ahrens and Kleinfeld (2004) noted the increase in vibrissal motor cortical LFP power just prior to exploratory whisking occurred between 4 – 10 Hz, which could partially explain the lower bound on the frequency range we saw in mouse C's bursting activity around reward availability (5-17 Hz). Previous studies reported 7-11 Hz coherence increases between rat barrel cortical LFPs and whisker EMG during goal-directed behavior, compared to exploratory whisking (Ganguly and Kleinfeld, 2004). Given the strong reciprocal connectivity between vM1 and vS1 (Farkas et al., 1999; Mao et al., 2011), there could be similar coherence or power increases in motor cortex during our goal-directed task. This increase in vM1 power or vM1-vS1 coherence potentially involves modulation of whisking or other orofacial behaviors, at higher frequencies due to our task requirements, even though much of the motor drive for orofacial behavior originates in the brain stem (Kleinfeld et al., 2014). We saw a large bilateral decrease in ~0 – 8 Hz LFP-EMG coherence for all mice that temporally coincides with the reward seeking response period (Figure 7C). We speculate this large decrease in coherence was due to the cessation of goal-directed whisking, and not due to decreased whisking alone because

the EMG power still increased over baseline in these frequencies at this time (Figure 7B). Cycle-by-cycle control occurs at lower frequencies than tonic "hold" beta, however, so only the low frequency EMG activity (up to the low beta range) may therefore be related to goal-directed whisking.

In sum, a tonic "hold" (or similar) motor output may have been the successful neuromodulation strategy responsible for the beta band range EMG activity and vM1-EMG relationships seen.

Increases of Beta Burst Durations

Lundqvist et al. (2016) reported no change in the duration of beta bursts, and claimed this as a predicted outcome of their model. The change in beta burst duration observed in the present study may result from the fact that beta band activity is directly tied to reward in our task, whereas features of beta band activity in Lundqvist et al. are correlative measures. In other words, our subjects had an interest in potentiating or activating neural connections involved in beta band activity—and received neurofeedback expressly for them to do so. If the increasing burst durations are a prolonging of normal oscillatory bursting processes, then it is possible that plasticity mechanisms, such as an increase in $GABA_A$ conductance, are responsible (Baker and Baker, 2002).

Otherwise, if these bursts are direct drivers of motor activity (as we suspect), then the mice may simply be extending those motor processes in a more deliberate manner for reward. In support of this latter possibility, Sherman et al. (2016) built on previous work

to model the potential mechanisms of beta bursts in somatosensory cortex (Jones et al., 2009; Sherman et al., 2016) and found that broad proximal drive of layer 2/3 and 5 pyramidal cells, coupled with distal inputs consisting of 10 Hz bursts lasting a beta period could be responsible. Whisker control on a cycle-by-cycle basis or tonic "hold" beta synchrony may require vM1 to receive the commensurate number of distal input bursts from higher cortical areas, or from S1 as part of normal closed loop sensorimotor control. In this possibility, the increase in beta burst durations we see could simply be a result of intentional increases in the duration of motor output and not due to plasticity.

Controllable Beta vs. Automatic Beta

Given broad freedom to solve the task using a strategy of their choosing, all three mice performed successful trials approximately as follows: (1) poke the snout into the trial initiation port, and (2) remain within the trial initiation port while continuously performing various orofacial behaviors, including whisking, licking, biting, and sniffing, until the reward tone sounded. Throughout this basic procedure, the vM1-EMG relationship—as quantified by task-related LFP-EMG coherence changes and MUA-EMG spike-field coherence—appeared highly similar among the three mice. No "action" was being decoded, as the mice had no requirement other than to increase the BPR value, yet the EMG patterns and vM1-EMG relationships were relatively consistent between the three mice.

One possibility may explain the similarities: the requirement to increase the BPR value may not be as generic as we thought—while there may be more than one

mechanism to generate beta rhythms, they may not all be volitionally controllable (Wood et al., 2014). Beta band activity that is possible to neuromodulate may be a more "intentional" beta that occurs within a narrow behavioral context. For example, motor command adjustment preparatory-beta (Torrecillos et al., 2015) and tonic-hold beta (Baker et al., 1997) may be more intentional activity or is otherwise more "controllable", whereas error salience/forward model updating rebound-beta (Torrecillos et al., 2015; Cao and Hu, 2016; Tan et al., 2016) may be a more automatic process associated with neural plasticity and motor learning. The consistent vM1-EMG relationships and BPR increases during high EMG suggest the likelihood that the modulated beta rhythm was of the tonic-hold variety. This directly implicates tonic hold of those muscles for which the vM1 cortical area controls, implying the inter-mouse behavioral similarity. If these bursts are direct drivers of motor activity, then the increasing burst durations we found may simply be the mice extending those motor processes in a more deliberate manner for reward. The rapidity with which mice learn the task further supports the idea that the mechanism of beta power production may have been a pre-existing capability exploited for reward. Future neuromodulation studies may wish to consider the context and mechanisms under which the particular rhythm is generated to prevent over-generalizing their conclusions (Jones, 2016).

Relationship to Previous Studies

Our burst durations are similar to previous findings, which found most bursts to last ~2 – 4 cycles (Figure 8). We included frequencies lower than the beta range in our

burst analysis, which we believe is necessary due to the power increases seen in the trial-averaged spectrograms. Feingold et al. (2015) did not include frequencies lower than 13 Hz in their analyses, despite the visible increase in power down to ~10 Hz. This exclusion appears based on the traditional definition of the beta rhythm as 13 – 30 Hz activity, which empirical data does not always support (Whittingstall and Logothetis, 2009).

*** Address motor imagery SMR studies, and why this study more explicitly leads to the conclusion that particular behaviors/contexts lead to beta modulation. Self-reporting vs. behavioral & neural measurements. Specifically asked to do something. No temporal measurements (i.e. when did imagery begin/end and beta occur).

Specific Implementation of Experiment

The electrode arrays spanned a larger area than vM1, so it is possible that neural activity was not entirely representative of vibrissal-related activity, and could have also included upper lip/muzzle, forelimb, eye/eye lid, neck, digit, wrist, or jaw-related activity (Pronichev and Lenkov, 1998; Matyas et al., 2010; Tennant et al., 2011; Xu et al., 2012). Indeed mice appeared to perform various orofacial behaviors such as licking, biting, and sniffing while they remained in the trial initiation port, attempting to increase the BPR value. Beta range activity is more widespread than higher frequency activity, however (Miller et al., 2009b), so the inclusion of some extra-whisker pad activity may not overly confound our interpretation.

During preliminary work, we attempted to run this experiment using decreases in 13-30 Hz band power as a reward condition, but real time calculations of this power band were generally at ‘floor’ and decreases were not possible. However, a “brain switch” based on beta ERS could be more promising than one based on beta ERD for several reasons. First, the time course of beta ERD in motor cortex is highly variable across task conditions (Stančák et al., 1997; Tzagarakis et al., 2010; Fujioka et al., 2012). Second, beta ERS is a more robust phenomenon than ERD during both actual and imagined movement (Bai et al., 2008; Pfurtscheller and Solis-Escalante, 2009; Pfurtscheller et al., 2010; Solis-Escalante et al., 2010; Zaepffel et al., 2013); in one study a paralyzed subject moved a virtual wheelchair through motor imagery beta ERS with no false positive detections (Leeb et al., 2007). Third, beta SMRs are more prominent during periods of increased sensorimotor activity than during rest (MacKay and Mendonça, 1995; Takahashi et al., 2011). Finally, down-regulation of neural activity appears to be more difficult to achieve than up-regulation (Fetz and Baker, 1973; Kobayashi et al., 2010; Rouse et al., 2013; Clancy et al., 2014; Ramot et al., 2016), which is in strong agreement with our preliminary data. This may be an issue specific to our hardware setup, however, and we would not rule out other methods of training mice to decrease 13-30 Hz band power.

The increase in burst density prior to trial initiation and the decrease in density beginning around reward availability just prior to trial termination argue against motion artifact as the major cause of burst events because the required action for both trial initiation and termination is extremely similar (i.e. breaking an infrared beam in almost

identical nose ports, save the water spout in the reward port). Therefore, if motion artifact was the cause of the beta burst, we would likely see similar increases in the burst rate at both trial initiation and during and after termination. Instead, the burst density near trial termination peaks just before reward availability, and then decreases over the following ~400 – 600 ms, which tightly corresponds to the reward response latency while mice are actively moving towards the reward port.

Notable similarities between mice A and B and their differences to mouse C could be due to the unique online filter used to calculate mouse C's BPR. Mouse C had a lower range of frequencies in which power increased around reward threshold for both its LFPs and EMGs, for instance. The particularities of the real time filters used could have inadvertently taught the mice to solve the neuromodulation task through different neural and behavioral strategies, although we showed LFP- and MUA-EMG relationships were reasonably similar.

Feasibility of Beta SMR Brain "Switch"

Whether beta range activity used in this context is a viable brain "switch" is yet to be settled. Beta band activity appears to consist of bursts that increased in duration and not in amplitude, therefore a beta "switch" based on duration might be more successful. Additionally, the timescale on which mice reached reward threshold after trial initiation suggests that motor BMIs could be "switched" in a matter of seconds. This would be a long time by healthy standards and is not as fast as previous studies determining the idle state from spiking data (Velliste et al., 2014), but is a reasonable solution if only LFP

data is available and the threshold of activity needs to be difficult enough to achieve to avoid false positives.

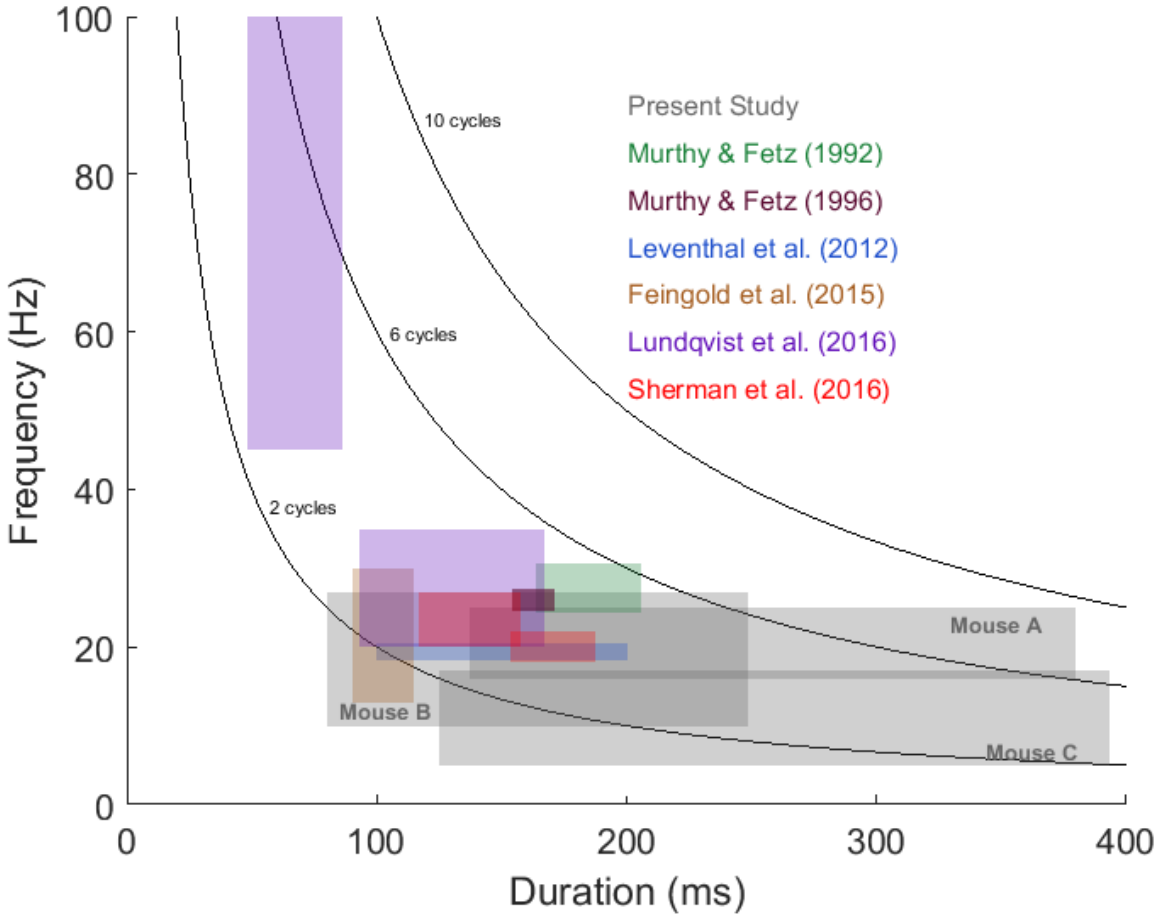

Figure 8: Burst results compared to previous descriptions

The present findings on burst frequency and durations are comparable to previous descriptions in the literature. Gray boxes: the 16th – 84th percentile (i.e. approximately the median ±1 STD) burst durations of the three mice from the present experiment. Colored boxes: reported frequency and duration ranges from previous studies (Murthy and Fetz, 1992, 1996; Leventhal et al., 2012; Feingold et al., 2015; Lundqvist et al.,

2016; Sherman et al., 2016). Black lines: the number of waveform cycles as calculated by frequency and duration (number of cycles as labeled).

CHAPTER THREE: Neural and behavioral strategies during adaptive decoding of a delayed saccade task

Introduction

Brain-machine interfaces (**BMI**s) are quickly becoming more usable (Collinger et al., 2013). A primary concern for BMIs, however, is the development of decoding algorithms that are able to adapt to changes in neural measurements that result from both endogenous and exogenous processes to the brain. Some factors are well known to affect neural activity and measurements, such as the level of attention, plasticity, alertness, gliosis, electrode drift, and alternative control strategies (Pichiorri et al., 2011; Collinger et al., 2013; Perge et al., 2013), but there are also less common factors like self-perceived ability to control technology and the presence of unusual neurotransmitters that can also have an effect (Witte et al., 2013). Changes in any one of these factors could lead to a decrease in BMI decoding performance. The goal of adaptive decoding is to allow appropriate changes to the decoding model so that BMI performance is maintained or improved, despite fundamental alterations to the underlying neural activity or measurements.

An issue related to neural alterations and measurement drift is that decoder performance can be dependent on daily calibration sessions, which delays implementation and often requires the involvement of an onsite expert. It would benefit

users to reduce or eliminate the amount of calibration required for optimal BMI decoding. In this study, we attempt to address both issues; can we reduce the number of calibration trials required before a BMI becomes usable, and can we also adapt the decoding algorithm to ongoing changes in neural activity?

One method of reducing calibration time would be to pre-train a decoding model and then use that model as a basis for iterative, closed loop learning (Williams et al., 2013). This training method would be one example of adaptive decoding. Due to the importance of using non-linear, spatiotemporal decoding methods for neural data (Chapin et al., 1999), we used a Jordan network, otherwise known as a simple recurrent neural network (**RNN**).

An additional method of reducing calibration time is to avoid using spiking data as the basis for decoding. Spike sorting is generally a manual process performed by experts that produces questionable results (Harris et al., 2000; Wood et al., 2004), and minor electrode displacements can drastically alter the recorded spike waveforms, negating any previous spike sorting (Harris et al., 2000). A potential alternative is to use power in the 80-500 Hz frequency band as a reasonable correlate of spiking activity (Manning et al., 2009; Miller et al., 2009a; Miller, 2010). Using this frequency band precludes the need for a sorting process and may be more resilient to changes in electrode placement, given the distributed nature of the signal sources. Furthermore, as with decoding algorithms that use multi-unit activity, there is less need for pre-session trials in which the researcher records data purely for the purpose of spike sorting (DiGiovanna et al., 2009).

Researchers have attempted closed-loop paradigms to improve the decoder, as well as improve the method of stimulation (Fernandez-Vargas et al., 2013).

The frequency of feedback in adaptive BMIs is an important factor in how subjects learn. Many previous closed loop adaptive decoders have been used in dynamic control tasks in which the subject receives continuous feedback (Chapin et al., 1999; Wessberg et al., 2000; Taylor et al., 2002; Wolpaw and McFarland, 2004; Gage et al., 2005; Vidaurre et al., 2006; Wahnoun et al., 2006; DiGiovanna et al., 2009; Vidaurre and Blankertz, 2010; Ludwig et al., 2011; Gürel and Mehring, 2012). Relatively few studies looked at adaptation in discrete classification tasks where the feedback is only given at the end of a trial after a decision has been made, either offline (Bishop et al., 2014) or in a closed loop (Bryan et al., 2013). Based on ongoing work (Brincat et al., 2013a, 2013b), we first used an adaptive model offline to classify upcoming saccades using 80-500 Hz delay period activity. Offline decoding success was high enough to warrant maintaining the same paradigm in a closed loop adaptation experiment.

Preliminary offline work of an adaptive simple recurrent artificial neural network (**RNN**) showed promising results that suggested closed loop decoding and adaptation of discrete eye movements was feasible (Torene et al., 2013). Performance of RNNs pre-trained on previous sessions began above chance (~50 – 65%; chance = 16.7%) and increased over the first ~100 trials to asymptotic performance (~70 – 80%), which was comparable to LDA performance which required 300 correctly decoded training trials at the beginning of each session (Brincat et al., 2013b). RNNs utilizing generalization techniques (see *Adaptation Strategy*) had higher next day performance than non-

generalized RNNs. Based on these results, we attempted to run the adaptive RNN in a closed loop with two monkeys, but the closed loop results did not match offline results, which is not unexpected (Koyama et al., 2010). We saw an initial increase in performance over the first few adaptive sessions in the two monkeys (from ~25% to ~75% correct for monkey C, and from ~36% to ~54% correct for monkey J), but performance then fell for the remaining five sessions (~23% for monkey C, and ~31% for monkey J).

Methods

Overview

This study utilized existing data from an ongoing experiment (Brincat et al., 2013a, 2013b). Two macaques performed a delayed saccade task during which neural activity was recorded from three cortical areas associated with eye movement planning and execution: the dorsolateral prefrontal cortex (**dlPFC**), frontal eye field (**FEF**), and supplementary eye field (**SEF**) (Bruce and Goldberg, 1985; Funahashi et al., 1989, 1993; Schall, 1991; Dias and Bruce, 1994). A simple recurrent neural network (**RNN**) used the power of the 80-500 HZ band as inputs to decode saccadic eye movements. The RNN was initially batch trained, and then iteratively trained on subsequent novel data.

Subjects

Two male monkeys were used in this study: a *macaca fascicularis* (monkey C), and a *macaca mulatta* (monkey J). Both monkeys were implanted in areas dlPFC, FEF,

and SEF with three 32-channel electrode arrays. Monkey C was implanted in the left hemisphere, and monkey J was implanted in the right hemisphere.

Behavioral Task

A 6-choice delayed saccade task was performed by monkeys C and J in two conditions: eye-control and brain-control (Figure 9). In the eye-control condition, crosshairs appeared in the center of a computer monitor, indicating the start of the trial. The subject was required to maintain fixation on the crosshairs for 500 ms and a target briefly flashed in one of six possible locations. After a delay period of 750 ms from the disappearance of the flashed target, the crosshairs disappeared and the macaque was then allowed to saccade to the location of the target. If the macaque made a saccade to the correct target, the trial was considered a success and a juice reward was provided. In the brain-control condition, neural activity recorded during the 750 ms delay period was used as input to the RNN (see *Data Collection and Decoding Features*) and reward was dependent on the output of the decoder.

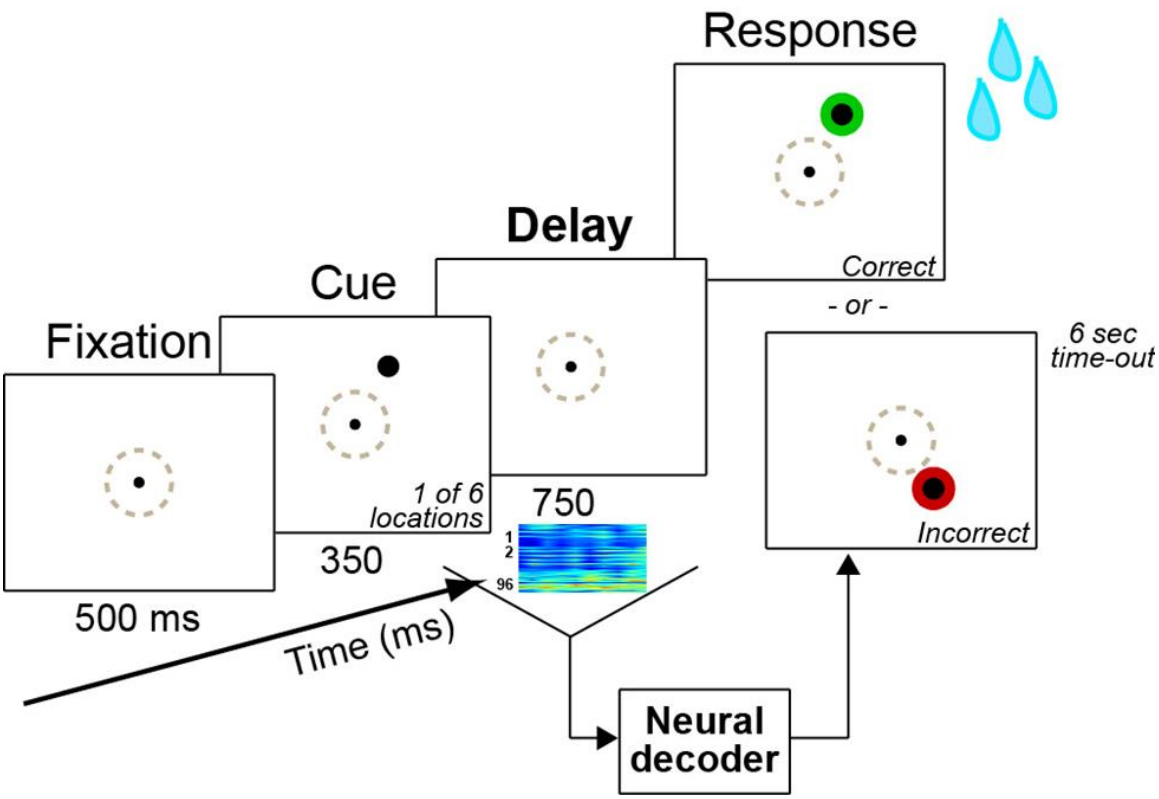

Figure 9: Delayed saccade task

If subjects maintain fixation on a center target for 500 ms, a target is cued for 350 ms. After the target disappears, the 750 ms delay period elapses, during which the subject must maintain fixation on the center of the screen. Once the center fixation target disappears, the subject is then free to saccade towards the cued target (eye control trials), or the decoder makes a decision about the subject's intended target (brain control trials). Saccades or decoder decisions towards the cued target are correct and rewarded (top right), while saccades or decoder decisions to other targets are incorrect and not rewarded (middle right). Figure adapted from (Brincat et al., 2013b).

Data Collection and Decoding Features

Three 32-channel BlackRock arrays (96 total channels; 400 µm spacing) were implanted in the left (monkey C) and right (monkey J) hemispheres, in areas dlPFC, SEF, and FEF, which are associated with working memory and eye movement planning and execution (Bruce and Goldberg, 1985; Funahashi et al., 1989, 1993; Schall, 1991; Dias and Bruce, 1994). Data was recorded at 30 kHz. The signal was bandpass filtered between 80-500 Hz using a 3rd order Butterworth filter and the 750 ms delay period was divided into ten 75 ms time bins. Each trial provided a 96×10 input array to the RNN, where the power of the 80-500 Hz band for each of the 96 channels at time step t was sequentially input at time step t. The 80-500 Hz frequency band is thought to be a proxy for nearby suprathreshold neural activity (Manning et al., 2009; Miller et al., 2009a; Ray and Maunsell, 2011) and a reasonable substitute for spikes, as unit sorting is not necessary for good decoding performance (Stark and Abeles, 2007; Fraser et al., 2009). Furthermore, this band was determined to have the highest decoding power in a LDA decoder from among the delta, theta, alpha, beta, and gamma bands, as well as from both MUA and single unit activity (Brincat et al., 2013a). Details of the LDA decoder are provided elsewhere (Brincat et al., 2013b). All data used in this study came from trials that were deemed correct by the closed loop LDA decoder.

All computation and analyses were from custom written scripts in MATLAB (R2011a, 64 bit). The computer on which the simulations were run had an Intel Core i5-2400 (3.10 GHz) CPU and 8 GB of RAM.

RNN Overview

We used a simple RNN to decode neural activity, with the intent of capturing any oscillatory dynamics occurring during the delay period (Elman, 1990). Recurrent connections in a simple RNN use the outputs from the previous time step as additional inputs within the same layer for the subsequent time step. The RNN consisted of one input layer with 97 nodes (96 electrodes + 1 bias), two logistic hidden layers with 111 and 61 units that were fully connected with a delay of one time step (110 hidden + 1 bias in the first layer and 60 hidden + 1 bias in the second), and one softmax output layer with 6 nodes (one for each target). Softmax output values are constrained between 0 and 1 to represent the probabilities of their respective targets being the desired choice of the monkey at each 75 ms time step in the 750 ms delay period. Initial weights were randomly drawn from the normal distribution ($\mu = 0, \sigma = 1$) and were adjusted with gradient descent backpropagation.

At each decoded time step, the target with the largest output is tabulated, and the final answer for each trial is a weighted mode of the ten responses. The weighted mode favors RNN output closer to the time of eye movement than stimulus presentation to reduce any effect early stimulus artifacts may have on decoding (Mohler et al., 1973; Bruce and Goldberg, 1985; Thompson et al., 2005; Tremblay et al., 2015). It does this by weighting the output value at time t approximately as $w \sim e^{0.1t}$, for $t = \{1 \leq t \leq 10,\ t \in \mathbb{N}\}$. For example, it is possible that the responses for the first six time steps will conclude that target A is the correct target and the responses for the last four time steps will

conclude target B is the correct target, which means the final output will be target B, rather than target A (i.e. $\sum_{t=1}^{6} e^{0.1t} \sim 8.6$ is less than $\sum_{t=7}^{10} e^{0.1t} \sim 9.4$).

Adaptation Strategy

An initial RNN model for each monkey was batch trained using correct trials from several previous static LDA decoding sessions and then iteratively trained after each sequential trial during the following adaptive decoding session (DiGiovanna et al., 2009; Vidaurre et al., 2011). The values of the learning and regularization parameters were decreased for online adaptation so that weight modifications would build slowly upon the batch-trained model (see Appendix for specific learning parameter values). Iterative gradient descent learning performed on an initially biomimetic decoder with a small enough learning rate should be relatively easy to learn by subjects (Danziger et al., 2009; Sadtler et al., 2014).

Previous work showed that lack of generalization is a potential issue with adaptive decoders (DiGiovanna et al., 2009), so in preliminary work (Torene et al., 2013) we looked at four regularization techniques to increase the initial performance of models trained from previous sessions' models: dropout (Srivastava et al., 2014), L2 weight decay (Krogh and Hertz, 1992a), artificial data generation using temporal bootstrapping (Barton and Schruben, 2001), and Bayesian targets (see Appendix for more details). Dropout approximates the training of 2^n different network architectures (where n is the total number of hidden units) which is a form of bagging for RNNs. L2 weight decay ensures that only meaningful weights stay large and that weights are less biased by the specific training set. Input resampling helps decrease the variance of the model by

artificially increasing the size of the training set. Bayesian targets relaxes the assumption that the neural activity is always representative of the true intention (or working memory) to saccade to the correct target; altering training data based on reasonable assumptions about user intent potentially contributes more performance increases than even retraining a decoder (Fan et al., 2014).

Several RNN decoding models were generated for each monkey prior to online closed loop adaptation. The best model for each monkey, as determined by test error rate during batch training, was selected as the starting point for adaptation. Preliminary work showed the generalization techniques reduced model bias and that models with higher batch performance had higher initial performance on subsequent sessions (Torene et al., 2013).

Offline Training Paradigm

We looked at two sets of three consecutive days of data, for a total of six days of data. Both sets of data were collected in August (15-17) and November (26-28) of 2012. In these sessions, the macaque's performance was determined by a static LDA decoder, although the macaque was also allowed to move its eyes toward the target. For this offline analysis, only trials that were correctly decoded by the LDA decoder were used to train and test the RNN.

RNNs were trained in two phases for this study (Figure 10). The first phase utilized batch training on data from four of the six days (August 15-16 and November 26-27; "previous days"). The second phase took the resulting models from the previous days and applied online learning to them using data from a separate set of four days (August

16-17 and November 27-28; “subsequent days”). Each model was trained on two consecutive days of data in order to look at the feasibility of using historical data as a basis for the decoding of future data. On the previous days, each model was batch trained on the entire day’s data, using 70% training, 20% validation, and 10% testing blocks chosen at random to find the optimal model for that day. Each batch iteration consisted of a block of 500 random training trials.

On the subsequent days, the batch trained model from the previous day was trained online using sequential trials, with 79% training and 20% validation blocks. Validation trials were selected randomly to avoid the possibility that we would be training towards late neural patterns of the monkey. Remaining trials in the training set, however, were presented in chronological order to simulate original conditions as best we could. The remaining 1% of trials was eliminated purely for programming expedience and is ignored for all offline analyses (these trials amounted to 17 and 11 trials per session for August and November, respectively, out of a total of ~1635 trials for each session in August and ~1068 trials for each session in November). Each online training iteration consisted of a cumulative block of all training trials up to the current trial.

Previous day batch model iterations were selected against the validation block, and final batch model performance was measured against the testing block. Test performance of the batch models was used in all final model selections. Next day online models were not chosen and therefore performance for each online training iteration was measured directly against the validation block, as that serves as the theoretical performance of the model for some block of unseen trials.

Batch training parameter values were determined through semi-automated exploration of performance for August 15. Once reasonable values were found manually, they were iterated upon to reach "best" parameter values for the number of hidden units and layers, learning rate, dropout rate, L2 weight decay, temporal input bootstrapping rate, and Bayesian target value. A three point moving average of test error was used to determine the best value for a given parameter, in which the center point with the lowest average test error was defined as the best value. Once found, the next parameter was iterated upon to find its best value. In this way, the cycle was repeated until lowest test errors were obtained from whole day batch training.

Appropriate online training parameter values were determined in a similar fashion as batch training parameter values, except that only August 15 was tested, and a relative proportion of batch training parameter values was chosen for all remaining online training parameter values. Many learning parameters were scaled down for online training so that learning built slowly upon the existing batch trained model. To determine whether the models would perform consistently, ten iterations of the three day training sequence (1st day batch/2nd day online, 2nd day batch/3rd day online) were run for both August and November, using both standard and generalized RNNs, for a total of 80 previous day batch/subsequent day online training simulation pairs.

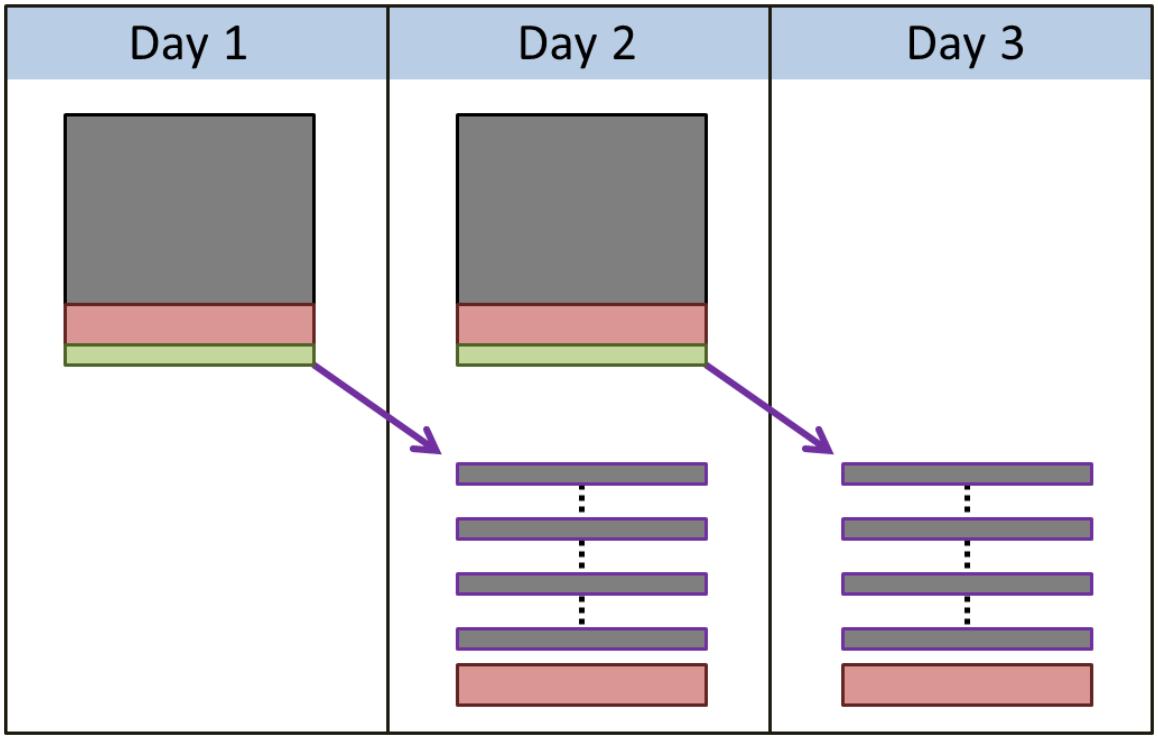

Figure 10: Offline model adaptation sequence

"Previous day" models (Day 1 and Day 2, top) are trained and selected through the standard training-validation-test (black, red, and green, respectively) data sets. Best models are selected via the highest validation performance, and the test performance is the final performance statistic considered. "Subsequent day" adaptive models (Day 2 and Day 3, bottom) are trained via the training data set, and tested against the validation data set after each new training trial. The running validation performance is the final performance statistic considered.

Closed Loop Training Paradigm

Closed loop training used the same batch training and online training parameter values as in the offline preliminary work. Monkeys C and J performed 5 and 6 (respectively) preliminary static LDA decoder training sessions and all correct trials from these sessions for each monkey were used to build 10 batch models, from which the best model was used as the basis for online training. At the beginning of each static LDA decoder session, monkeys performed 300 training trials in which reward was contingent on eye movements and not on decoder output. There were no transition sessions during which there was a gradual decrease in the number of training trials from 300 to 0 (i.e. immediate adaptive decoding). Each monkey performed 9 adaptive RNN decoding sessions. After each session was completed, the correct trials from the session were added to the pool of initial batch decoder training sessions to create a new initial model for the next session. Ten adaptive sessions were planned, but the experiment was stopped early due to fears that low closed loop task performance would be detrimental to future task performance in unrelated experiments.

Reconstruction of 80-475 Hz Power

We reconstructed 80-475 Hz power in each of the 96 electrodes during the delay period using the power of low frequency activity. The power of various frequency bands was calculated by bandpassing the entire session's offline LFP using a non-causal 6^{th} order Butterworth filter between the two frequencies of interest, then taking the RMS of the Hilbert envelope during the ten 75 ms time steps of each delay period. Phases were

also calculated during exploration, but phase information was not used in the final iteration of this process, so we will not go into detail.

A feedforward artificial neural network (**ANN**) with one hidden layer of 100 logistic units was used for 80-475 Hz power reconstruction. The ANN was regularized using dropout, and L1 and L2 weight decay.

Results

Offline Performance

Ten iterations of the three day training sequence (1st day batch/2nd day online, 2nd day batch/3rd day online) were run for both August and November, using both standard and generalized RNNs for a total of 80 previous day batch/subsequent day online training simulation pairs.

Generalization of the RNN is the key to improved performance. Figure 11 shows a representative previous day batch training sequence comparison between a standard RNN with no generalization features, and a generalized RNN with dropout, L2 weight decay, artificial input bootstrap generation, and Bayesian targets. Features of this example that are common to many other examples is that the standard RNN reaches its maximum validation value earlier than the generalized RNN, but that the training and validation rates diverge noticeably—the training success rate often reaches 100%, but the maximum validation rate averages 78% and 70% for August and November, respectively. This is in contrast to the generalized RNN, in which the training and validation rates remain highly similar throughout training, and the maximum validation rate averages

83% and 80% for August and November, respectively. Generalization also increases batch test performance (Figure 12).

In all four online training cases (August 16 and 17, and November 27 and 28), the generalized RNNs outperformed the standard RNNs (Figure 13). In all but one instance, there was a significant difference between initial online performance of the generalized and standard RNNs. Final online performance of the generalized RNNs was significantly higher than the standard RNNs for all four days ($p < 0.01$ for each case, Holm-Bonferroni corrected).

To understand whether it is possible to estimate subsequent online performance of the generalized RNNs based on previous day batch performance, we calculated Pearson's correlations between three performance measures: 1) previous day batch test performance and subsequent day initial online validation performance; 2) previous day batch test performance and subsequent day final online validation performance; and, 3) subsequent day initial online validation performance and subsequent day final online validation performance (Figure 14). All three correlation measures were found to be significant ($\rho = 0.34$, $p < 0.05$; $\rho = 0.45$, $p < 0.01$; and $\rho = 0.72$, $p < 0.00001$, respectively; Holm-Bonferroni corrected). These correlations suggest that best practice is to generate several models on the previous day, and use the model with the highest performance and ignore those with poor performance to increase the chances of achieving good success rates on the following day.

Final online performance of the generalized RNNs was significantly higher than initial online performance ($p < 0.01$ for all subsequent days, Student's t-test, Holm-

Bonferroni corrected). Mean initial online performance of the generalized RNNs was 61.2% correct, while mean final online performance was 77.4% correct. Individual days had better or worse performance.

Decoding performance of individual targets for the previous day batch trained models suggests performance suffers from unilateral implants. The upper and right portions of visual space were better decoded than the lower left portion of visual space (Figure 15). This is likely due to the left-lateralized implants, as left dlPFC, FEF, and SEF are thought to be responsible for eye movements to the right half of visual space.

We looked at the largest 1% of weights (by magnitude) in each layer for a random model for Monkey C to determine whether any general statements could be made regarding the appropriateness of using RNNs as decoders (Figure 16). Three key features suggest RNNs to be highly appropriate decoding models. First, most of the largest weights from the input layer are fed forward from SEF, which indicates it to be the most informative cortical area for target intention. This was shown independently through related work in the same monkey (Brincat et al., 2013b). Second, the largest weights in the output layer were feedforward connections from the 2nd hidden layer to targets 1, 2, and 6, which are the three contralateral targets to the implanted hemisphere. This is consistent with cross-callosal processing and is recognized as most informative by the RNN. Third, recurrent connections were major constituents of the largest weights in each layer, suggesting the importance of a recurrent model for decoding delay period activity instead of a simple feedforward model.

Finally, we also established the practical matter of whether closed loop decoding using the generalized RNN could be accomplished in real time. While it can take several hours to derive a base model from batch training, online training and decoding occurred faster than real time. We reran 10 iterations of the online training while timing them. It took an average of 149 ms (± 33 ms STD) per 750 ms delay to train the models online. There are ten time steps decoded during each delay (75 ms each), which indicates that each training iteration took approximately 15 ms. This suggests that models using the current parameters and computer hardware can be trained online at a rate of up to 67 times per second, ignoring the time required by preprocessing functions.

Altogether, we established that the generalized RNN should be used instead of the standard version, RNNs were a reasonable method for closed loop decoding of saccades, and they could be used in real time.

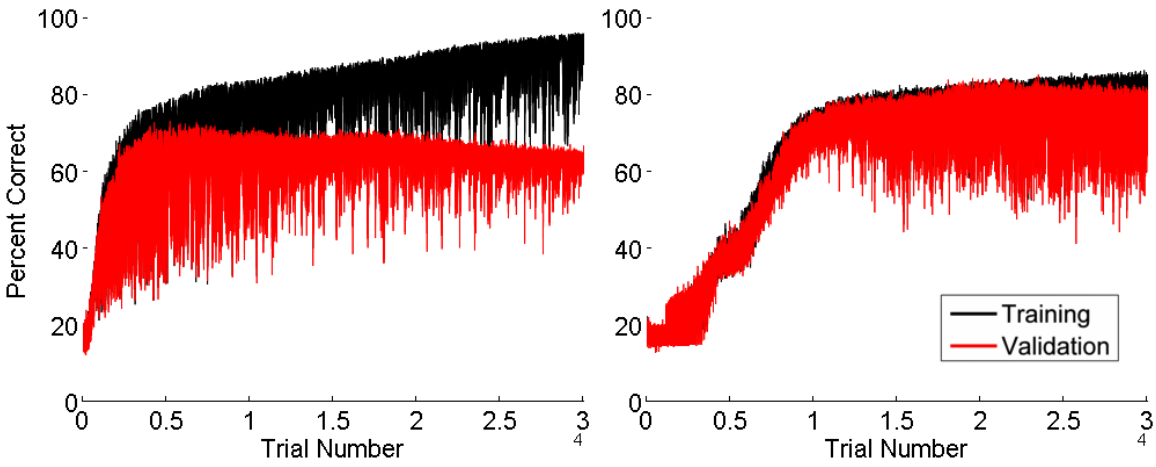

Figure 11: RNN generalization improves batch model validation rate

Generalization techniques improve offline batch model validation set performance. Left: training set performance of standard RNN approaches

100%, while validation set performance reaches a maximum of ~70%. Right: training and validation set performances of the generalized RNN both approach ~80%.

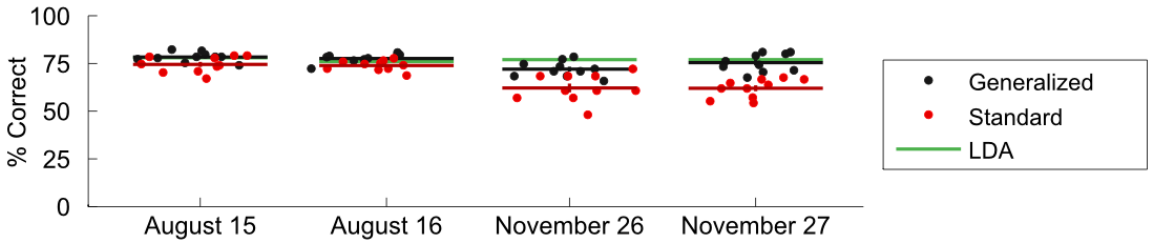

Figure 12: RNN generalization improves batch model test rate

Generalization techniques improve initial offline batch model test set performance. Shown are the test performances for the four "previous" offline decoding days. Ten model instances for each of the standard and generalized RNN show clear performance improvements of the generalized RNNs over standard RNNs for November. Closed loop LDA decoding performance is shown in green, for comparison.

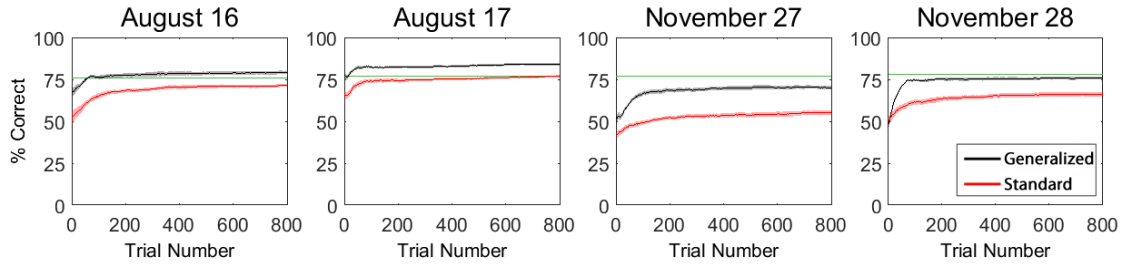

Figure 13: Online adaptation performance increases in offline simulations

"Next" day validation set performance shows online decoder adaptation improvement. Shown are the offline adaptation validation set performance

averages (± s.e.m.) of the 10 generalized models (black) and 10 standard models (red) on the four "next" days. Average generalized performance is comparable to static LDA performance, and indeed exceeds LDA performance for Aug 16 and 17. Note that validation set performance begins far above chance (~16.7%) at the beginning of each session, and increases to asymptotic performance within ~100 trials.

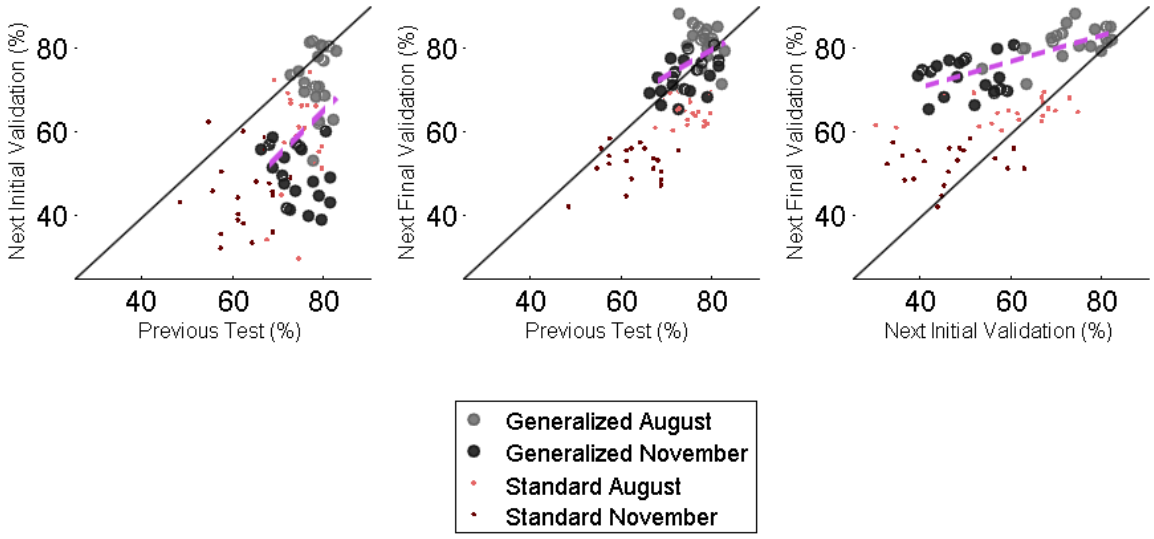

Figure 14: Better performing batch models have better online performance

Positive correlation of validation and test set performance statistics predict benefits for choosing models with the highest "previous" day performance. Left: a significant positive correlation was found between previous day test set performance and next day initial validation set performance ($\boldsymbol{\rho = 0.34}$, $\boldsymbol{p < 0.05}$; Holm-Bonferroni corrected). Middle: a significant positive correlation was found between previous day test set performance and next day final validation set performance

($\boldsymbol{\rho = 0.45}$, $\boldsymbol{p < 0.01}$; Holm-Bonferroni corrected). Right: a significant positive correlation was found between next day initial and final validation set performance ($\boldsymbol{\rho = 0.72}$, $\boldsymbol{p < 0.00001}$; Holm-Bonferroni corrected). These correlations make the strong case that the best performing "previous" day batch-trained models should be selected for "next" day online adaptation. Gray and black circles: generalized RNN performances. Red and pink dots: standard RNN performances for comparison (no statistical tests were performed on these data).

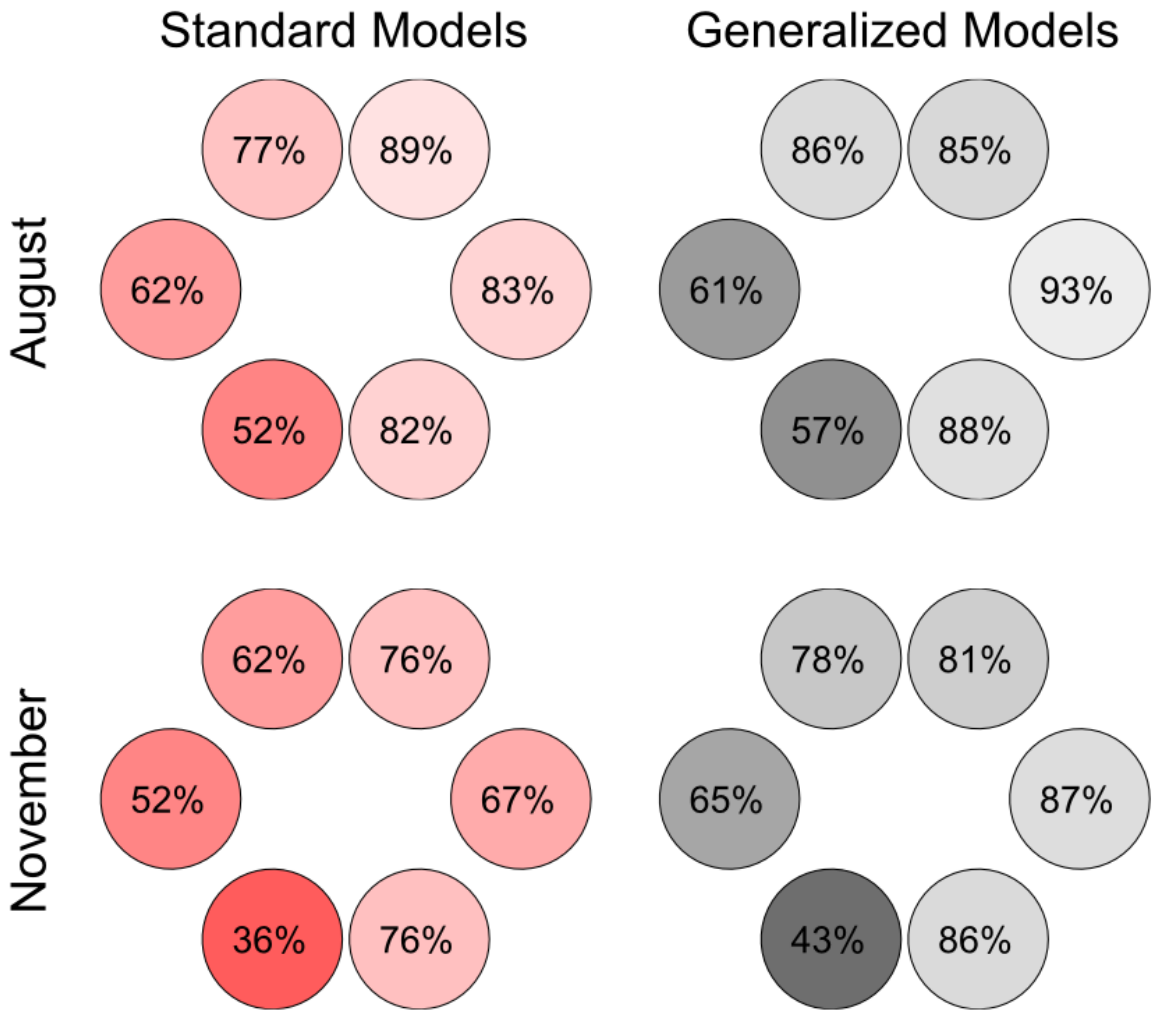

Figure 15: Higher contralateral than ipsilateral decoding performance

Target-specific batch test set performance is consistent with contralateral implant location. Monkey C had left lateral implants, and the targets in the right field of view (targets 1, 2, and 6) had higher decoding performance, which is consistent with previous studies and suggests RNN models learned relevant features of the data. Left column: standard RNN model performance shows that right lateral targets can have extremely high performance (e.g. 89% for target 2 in August), but left lateral target

performance suffers (e.g. 36% for target 5 in November). Right column: target-specific performance of generalized RNN models is higher than standard RNN models in all cases except two (targets 2 and 4 in August). These data suggest that bilateral implants would improve performance.

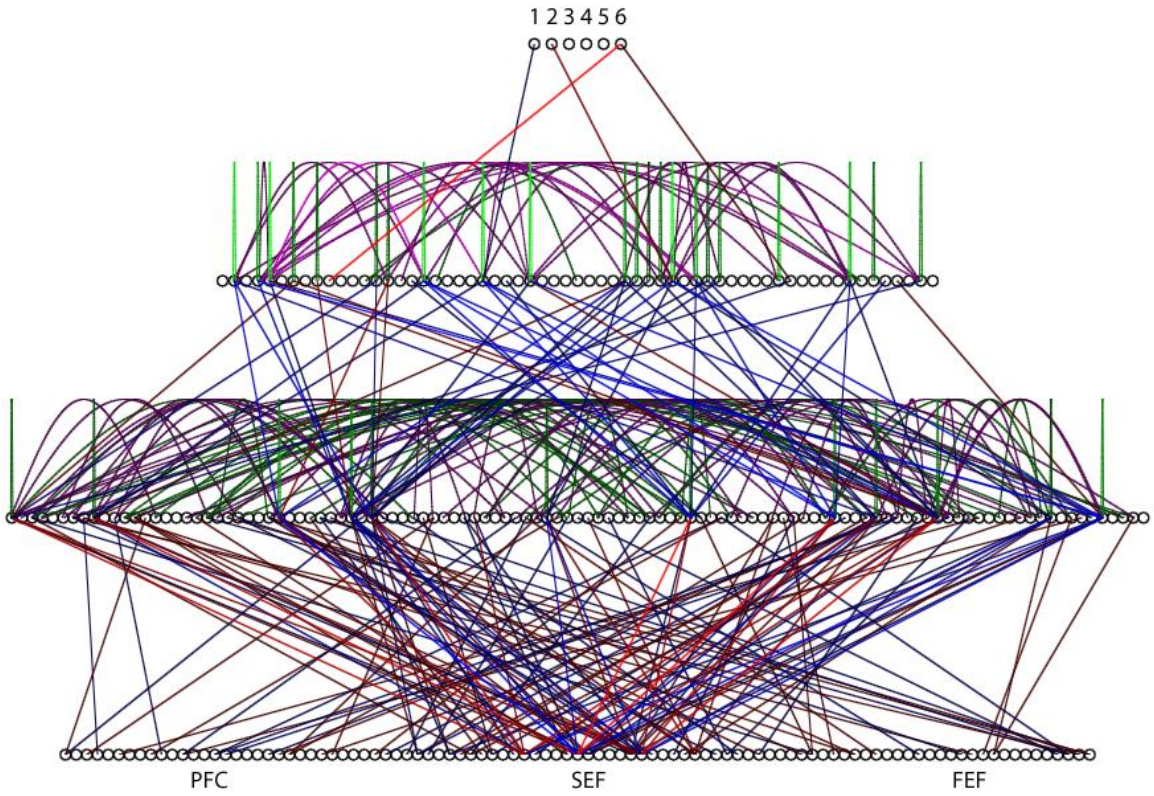

Figure 16: RNN weights consistent with related findings

Largest 1% of weights per layer, by magnitude, of one example model for Monkey C show three key features suggesting RNNs are reasonable decoding models: (1) most of the largest weights in layer 1 are derived from SEF electrodes, which independently confirms results from related work showing SEF to be most informative of target selection in monkey C (Jia dissertation); (2) many of the largest weights are recurrent instead of

feedforward, suggesting the importance of using recurrence; (3) the three targets in the output layer with the largest weights are contralateral targets to the implanted hemisphere (targets 1, 2, and 6). Red and green lines: positive feedforward and recurrent connections, respectively. Blue and purple lines: negative feedforward and recurrent connections, respectively.

Closed Loop Decoding Performance

Monkeys C and J each performed 9 closed loop adaptive RNN sessions after performing 5 and 6 closed loop static LDA decoding sessions, respectively. Adaptive RNN sessions began immediately under brain control, which contrasts with the static LDA decoding session paradigm of starting with 300 training trials performed under eye control.

Closed loop adaptive RNN performance was not as high as offline performance, nor was it as high as closed loop static LDA decoder performance (Figure 17). As may be expected from the offline decoding results, higher overall decoding performance was achieved for contralateral targets than ipsilateral targets (Figure 18). This trend was also seen within sessions (Figure 19).

The decrease in closed loop performance compared to the preliminary offline analysis could possibly be due to excessively large learning parameters, which would cause divergent weight changes. This possibility is unlikely due to the preliminary offline results, which indicate that performance should improve if the subject performs the task in complete ignorance of the decoder or changes to the decoder. The decrease in closed

loop adaptive RNN performance compared to the closed loop static LDA decoding performance is possibly due to the inappropriateness of a RNN for this decoding task. Again, this is unlikely due to excellent offline results.

An issue which could explain both unfavorable performance comparisons is that the task paradigm change was unexpected and aversive to the monkeys. Prior to the closed loop adaptive RNN decoding sessions, monkeys performed 300-500 training trials at the beginning of each session in which eye movements determined whether reward was received. Because healthy motor behavior is much more accurate than BMI decoders, it is possible that the monkeys were used to easy reward at the beginning and became frustrated with the adaptive decoding task, or possibly did not understand the new task parameters, yet understood the task parameters to have changed.

Session-level analysis suggests that most of the adaptive RNN performance was due to the performance of previous day batch-trained models and not due to within-session adaptation or learning for either monkey C (Figure 20) or monkey J (Figure 21). Least means square fits of session performance show that few decoders increased in performance within a session. Given that performance remained relatively constant within sessions, this suggests that poor closed loop performance was not due to problems with the model or learning parameters, but perhaps to factors external to the model such as monkey reaction to the change in task paradigm (i.e. starting sessions with brain control instead of eye control).

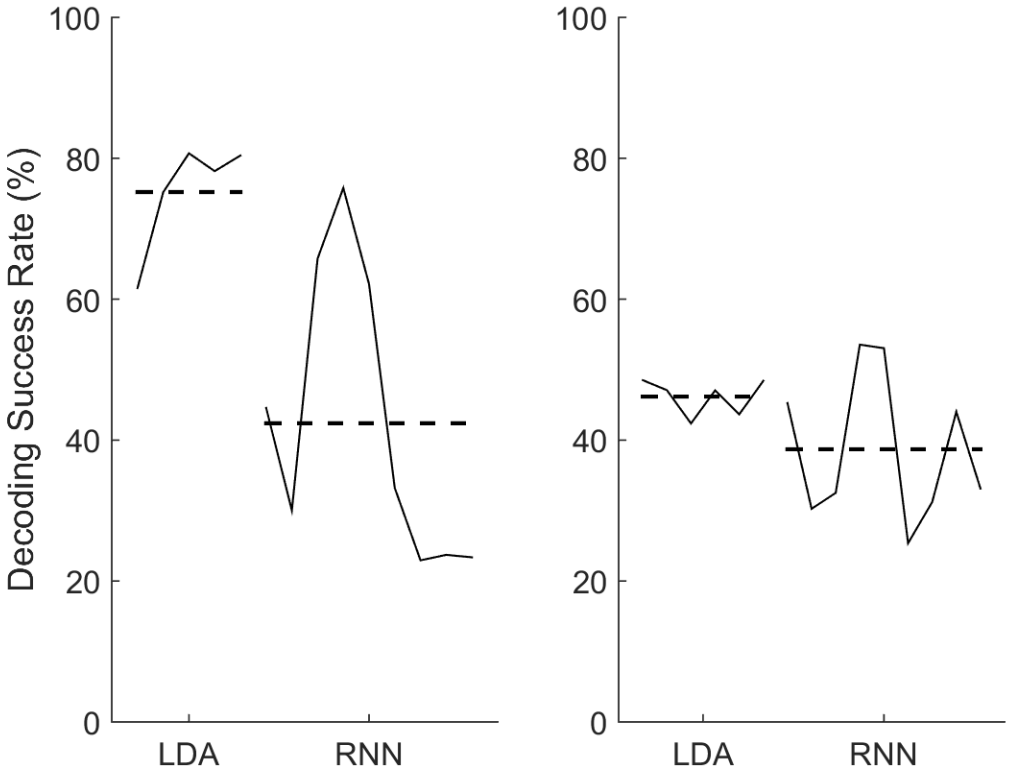

Figure 17: Closed loop RNN performance worse than offline and closed loop LDA

Closed loop adaptive RNN performance is more variable and lower performing than the static LDA decoder. For both monkeys, RNN performance was always above chance (16.7%), but did not match overall LDA performance. RNN performance for individual sessions occasionally exceeded mean LDA performance (RNN session 4 for Monkey C, and RNN sessions 4 and 5 for Monkey J), but this was the exception. Left: Monkey C static LDA and adaptive RNN performance. Right: Monkey J static LDA and adaptive RNN performance. Decoding success is determined by trials in which the decoders guessed the target (i.e.

monkeys successfully initiated a trial and did not prematurely break fixation).

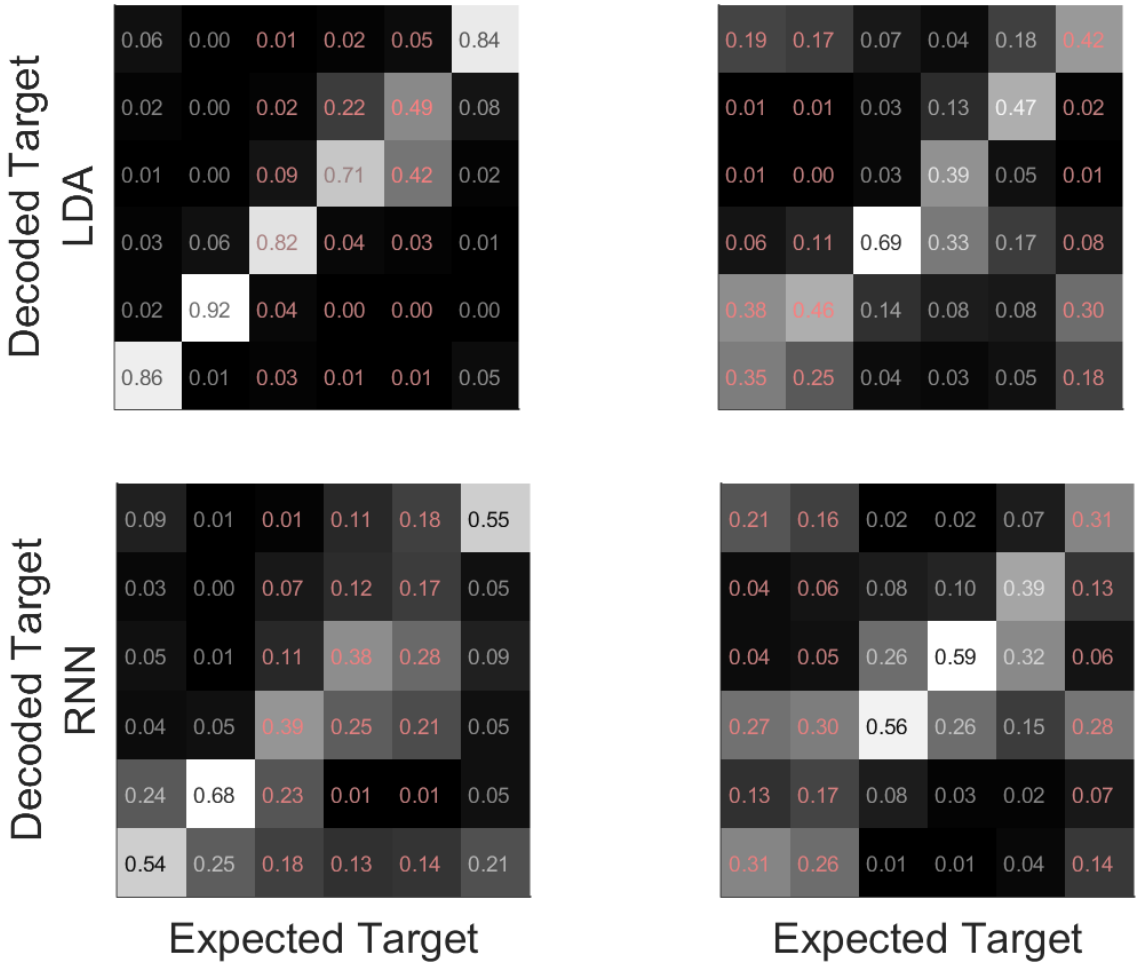

Figure 18: Higher contralateral than ipsilateral decoding performance

Closed loop adaptive RNN confusion matrix shows wider target confusion than static LDA decoder. Left column: Monkey C confusion matrices for LDA (top left) and RNN decoders (bottom left). Right column: Monkey J confusion matrices for LDA (top right) and RNN decoders (bottom right).

Grayscale numbering within the boxes indicates the Expected Target is contralateral to the implant; red numbering indicates ipsilateral.

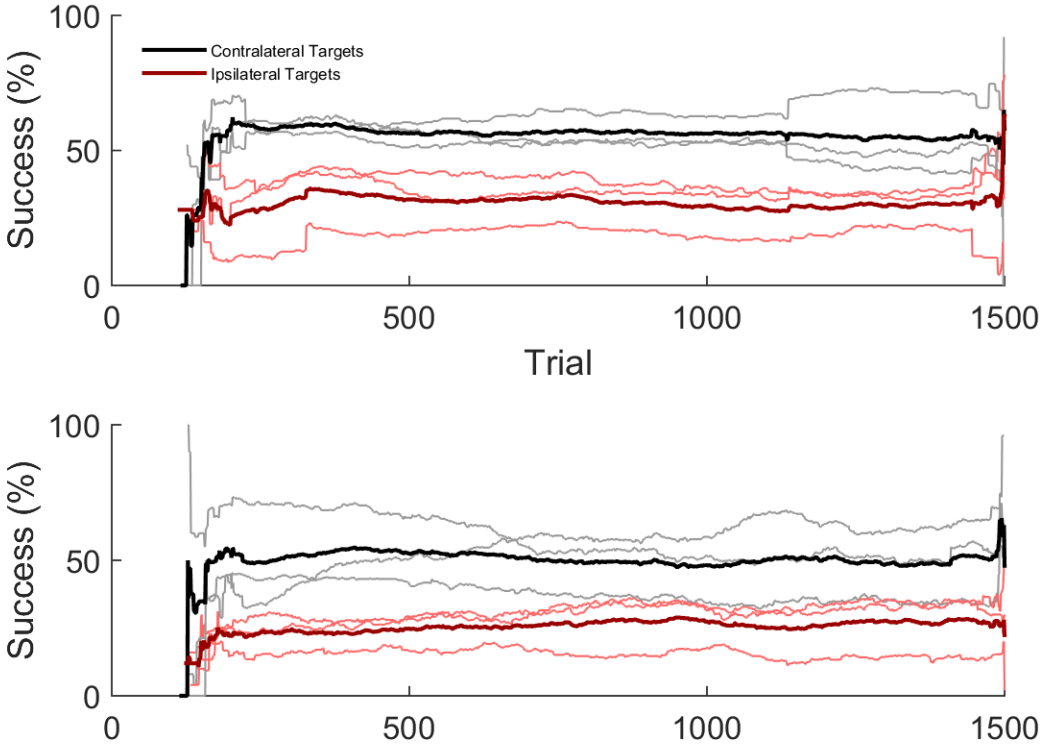

Figure 19: Higher online contralateral than ipsilateral performance

Contralateral closed loop RNN decoding performance is higher than ipsilateral performance and suggests the RNN models decoded reasonable features in the neural data. Top: Average contralateral (black) and ipsilateral (red) target performance for monkey C. Bottom: Average contralateral (black) and ipsilateral (red) target performance for monkey J. Note that large transients occur at the start of the data, due to peculiarities in individual session data and are averaged out as more sessions' trials

begin to apply to the data. These transients do not genuinely represent sudden changes in adaptation performance and long term trends should be paid more attention. Each trace is an average of 25-trial moving averages for each target, which is why success rates do not appear until ~125 trials.

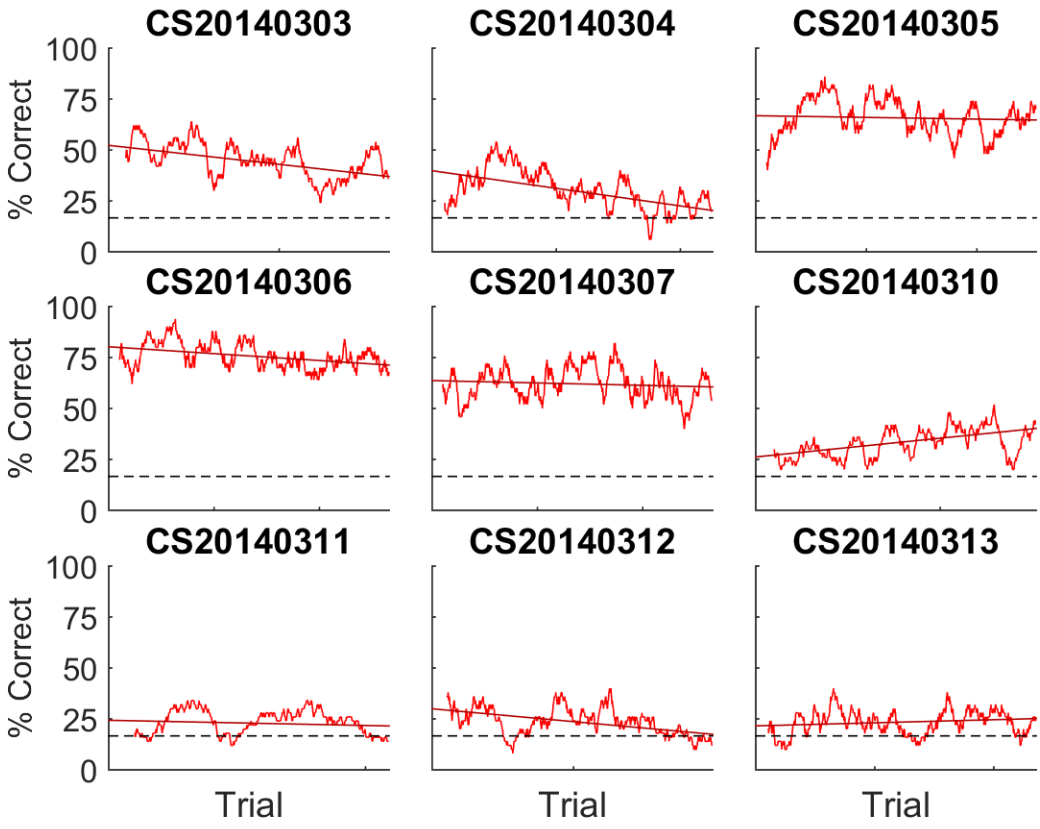

Figure 20: Monkey C performance not due to within-session adaptation

Few sessions show clear performance improvements, but performance generally begins above chance. This suggests that closed loop performance largely relied on a reasonable starting model. Red traces: moving average of 50 trials. Straight red lines: least mean square session performance. Black dashed line: chance performance (~16.7%). Black

titles: session identifiers. Sessions are ordered from left to right, top to bottom. % Correct is determined by trials in which the decoders guessed the target (i.e. monkeys successfully initiated a trial and did not prematurely break fixation).

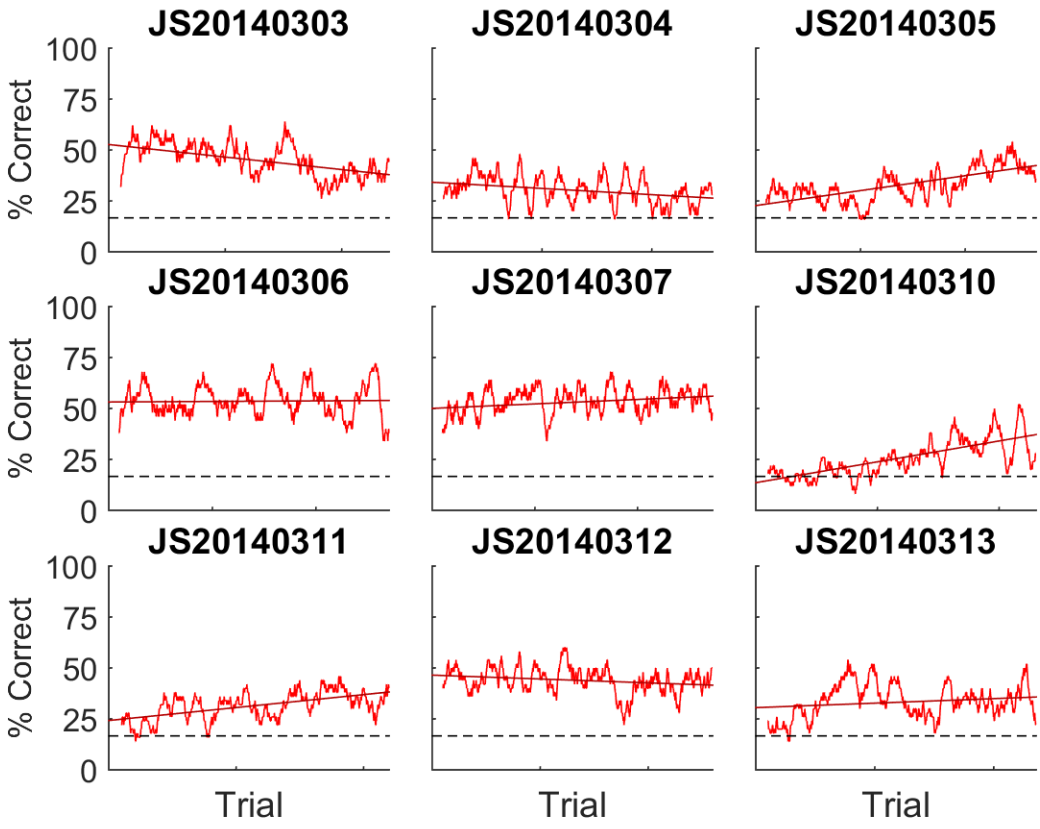

Figure 21: Monkey J performance not due to within-session adaptation

Few sessions show clear performance improvements, but performance generally begins above chance. This suggests that closed loop performance largely relied on a reasonable starting model. Red traces: moving average of 50 trials. Straight red lines: least mean square session performance. Black dashed line: chance performance (~16.7%). Black

titles: session identifiers. Sessions are ordered from left to right, top to bottom. % Correct is determined by trials in which the decoders guessed the target (i.e. monkeys successfully initiated a trial and did not prematurely break fixation).

Neural Correlates of Closed Loop Performance

At the end of the first week of adaptive RNN decoding, monkey C's FEF activity patterns during correctly decoded targets had changed substantially (Figure 22). On the contrary, monkey J's neural activity patterns appear to have remained relatively stable (Figure 23). Whether the changes in monkey C's neural patterns were led by attempted learning of the decoder, or were shaped by the decoder's adaptation is unclear. Expanding on this, it is unclear whether stable neural activity is a better strategy for decoding success, or if correctly decoded targets require particular activation patterns and the trials shown is merely a selection bias.

If initial models are generated based on data from previous sessions, then we wondered whether neural activity correlations between sessions had an effect on decoding performance of the current session (Figure 24). Monkey C's inter-session neural correlations were lower overall than monkey J, although monkey C's adaptive performance was often higher than monkey J. Nonetheless, we found a significant positive relationship between inter-session neural correlations and performance for each monkey (Figure 25; $p = 0.017$ and $p = 0.022$ for monkey C and J, respectively, Spearman's rank correlation, Bonferroni corrected). Therefore, consistent neural activity appears to benefit decoder performance, which questions why monkey C's neural activity

changed from the beginning to the end of the first week. Either the decoder forced monkey C's patterns to change by adapting improperly, or monkey C attempted to improve performance improperly, or changes in FEF activity, specifically, did not adversely impact the decoder.

Given the relative similarity of the cortical activation map between targets (Figure 22 and Figure 23), and yet the performance difference between ipsilateral and contralateral targets (Figure 19), we wondered how decoding performance within each target group changed over the course of an adaptive session. Muscle synergies are controlled by primary motor cortical neurons (Holdefer and Miller, 2002), so it is possible that proximal oculomotor areas are cooperative with respective to similar eye movements and/or movement goals. If so, then we might expect performance within contralateral targets and within ipsilateral targets to be positively correlated. On the other hand, if neural control of saccades or intended saccades within local cortical areas is competitive, we might expect laterality-specific performance changes across a session to be negatively correlated. We calculated within-session performance correlation of all target pairs and found the only ostensibly consistent pattern to be that ipsilateral targets increase or decrease in performance over time together (Figure 26). This could possibly suggest synergistic activity within working memory areas and oculomotor areas for ipsilateral goals. Alternatively, there could be such a relative absence of ipsilateral target processing that any increase in ipsilateral target decoding success "spills over" to nearby targets and increases the odds that a nearby target is decoded.

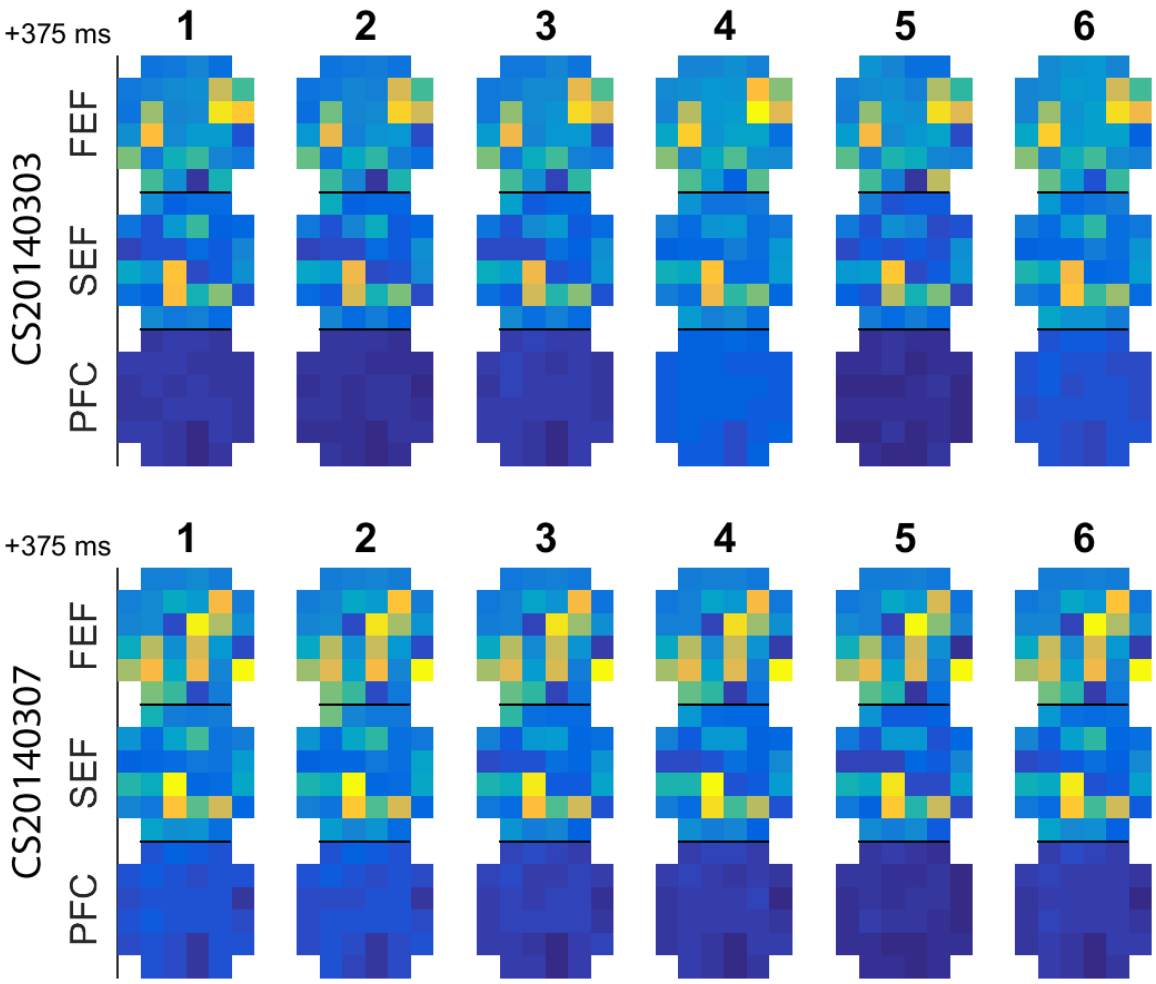

Figure 22: Monkey C neural activity patterns change

The neural activity patterns of monkey C changed over the course of the first week. Shown are the mean array activity patterns for the first (top) and last (bottom) session of the first week of adaptive RNN decoding during the 5th time step of CORRECT trials only. Activity in SEF remains substantially similar, but FEF activity changes considerably. It is unclear whether the neural pattern changes were due to the subject attempting to learn the decoder, and/or if preferred activity patterns were shaped by the

decoder adaptations. Shown are the six targets (as labeled). Note how similar the activity is between targets within the same session. All activity is normalized across all electrodes within a session. Although there does not appear to be much activity in PFC, there may still be substantial information decoded from this area.

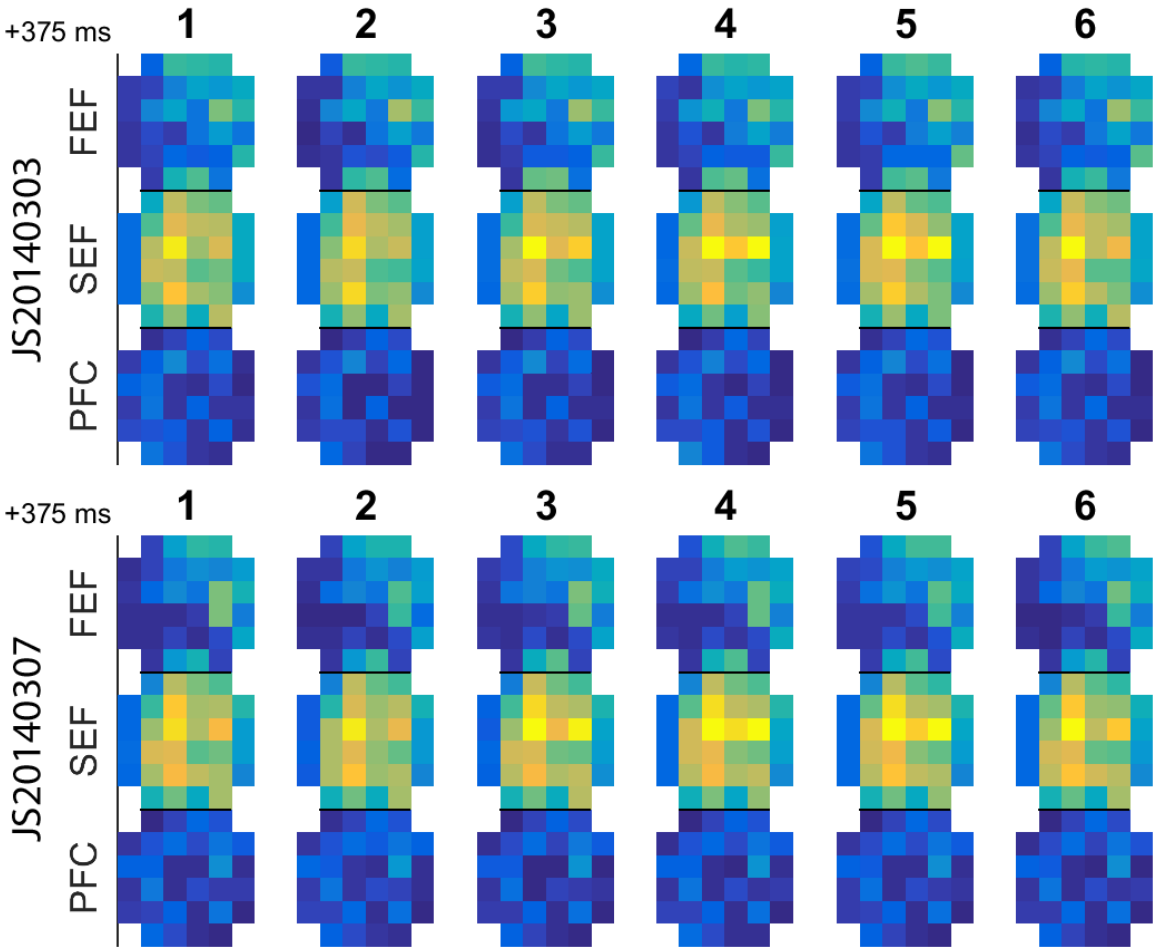

Figure 23: Monkey J neural activity patterns are stable

The neural activity patterns of monkey J appeared to change very little over the course of the first week. Shown are the mean array activity patterns for the first (top) and last (bottom) session of the first week of

adaptive RNN decoding during the 5th time step of CORRECT trials only. Activity in all three cortical areas appears to have remained substantially similar. Shown are the six targets (as labeled). Note how similar the activity is between targets within the same session. All activity is normalized across all electrodes within a session.

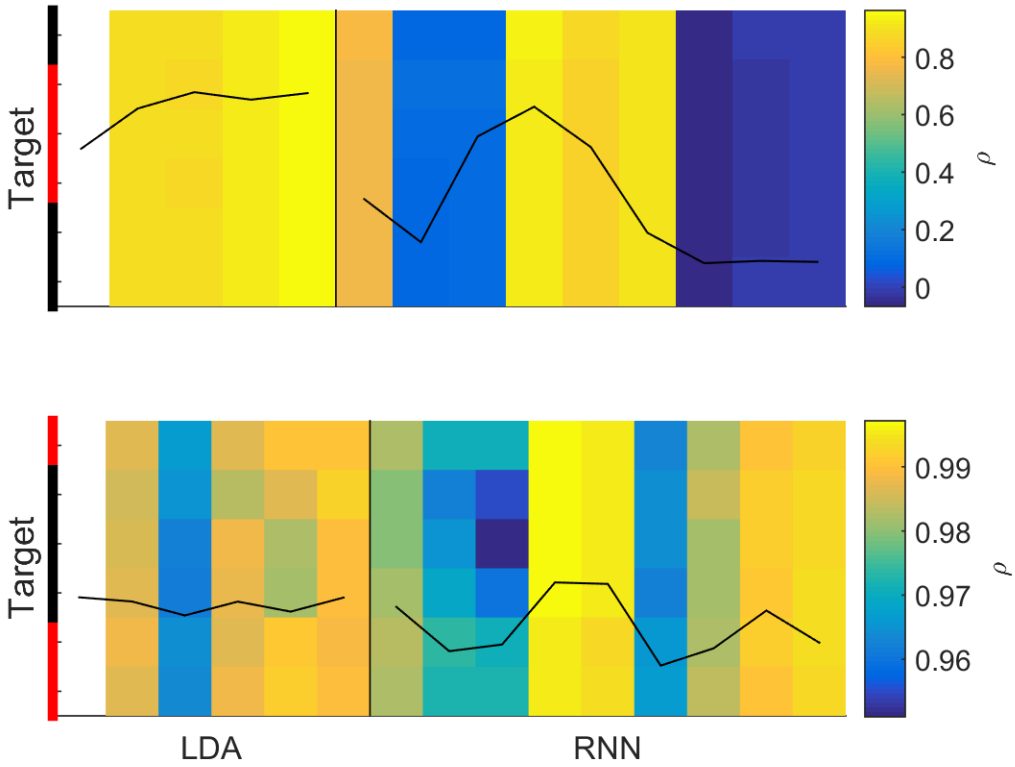

Figure 24: Performance related to neural similarity to previous sessions

Correlations of 80-500 Hz electrode power in "next" sessions with "previous" sessions suggest decoding performance is dependent on neural activity similarity to that which was used to construct the initial RNN model. Top plot: monkey C previous and current session correlation of 80-

500 Hz power. Bottom plot: monkey J previous and current session correlation of 80-500 Hz power. Black and red bars on left: contralateral and ipsilateral targets, respectively. Black line overlay: decoding performance (corresponding to the data from Fig. X), with top of graph indicating 100% success and the bottom indicating 0%.

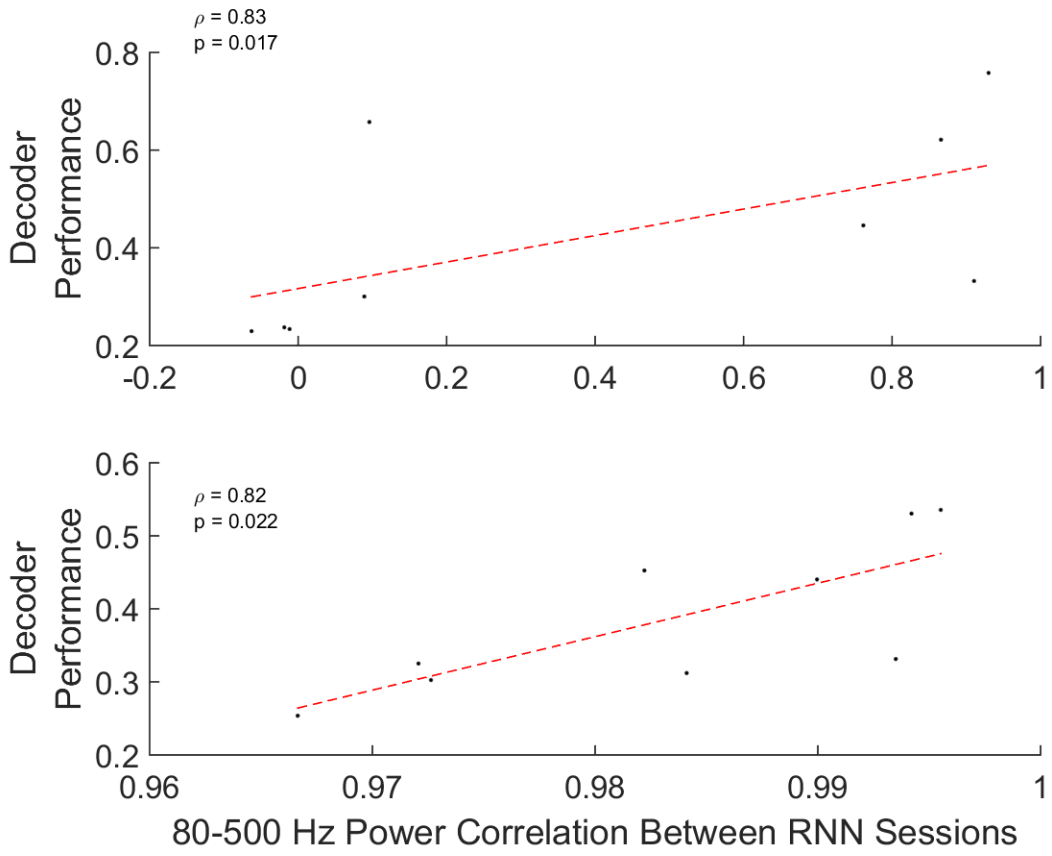

Figure 25: Better performance with consistent inter-session neural activity

Correlations of 80-500 Hz power during closed loop adaptive RNN "next" sessions to "previous" sessions' 80-500 Hz power show significant correlations to decoding performance for both monkeys ($p = 0.017$,

$p = 0.022$, Spearman's rank correlation, Bonferroni corrected). Top: monkey C. Bottom: monkey J. Red dashed lines: least means square fit.

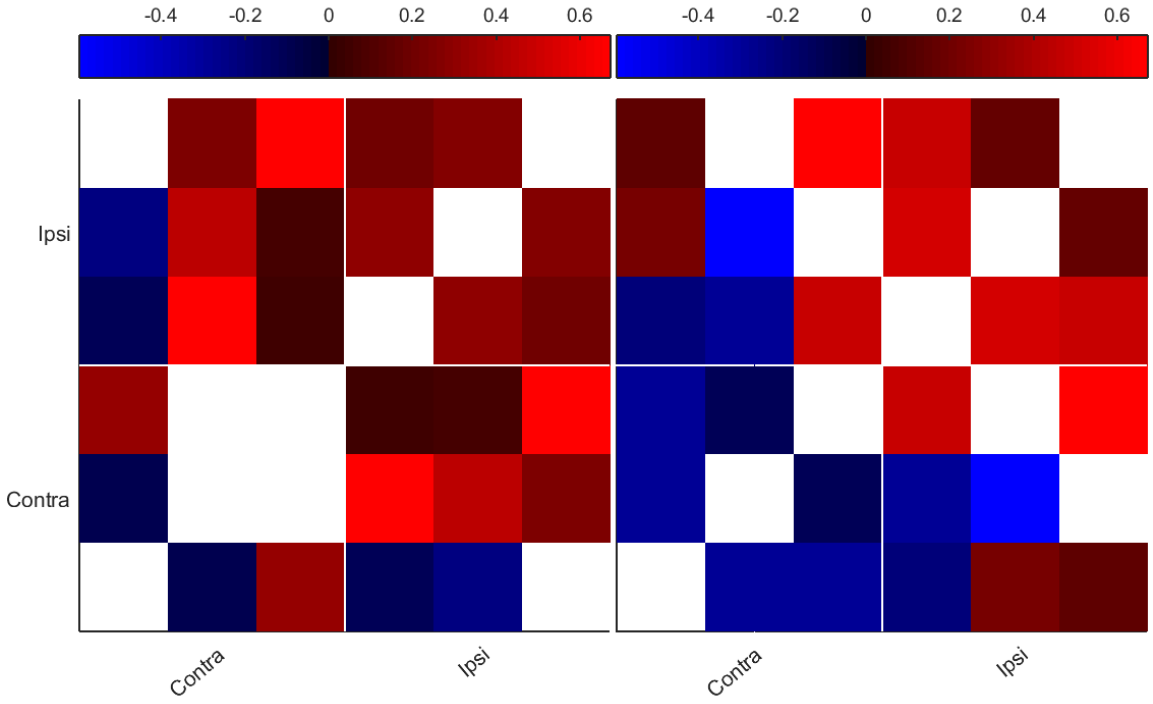

Figure 26: Decoding performance of ipsilateral targets changes together

Significant target performance trend correlations suggest ipsilateral target performances are positively correlated and may increase and decrease at similar times. All other correlations suggest no consistent pattern of inter-target adaptation or learning. Some targets conflict with each other during learning/adaptation, and some improve with each other. Shades of red and blue: significant positive and negative (respectively) target performance correlations across all closed loop RNN decoding sessions. The white diagonal squares from lower left to upper right are the same target correlations and are not shown. Contra: contralateral targets. Ipsi:

ipsilateral targets. Left: monkey C. Right: monkey J. Spearman rank correlations were Bonferroni corrected for the number of sessions (n=9) and target combinations (n=15).

Behavioral Correlates of Closed Loop Performance

We attempted to understand the behaviors associated with closed loop decoder performance. To differentiate which behaviors were due to trial type (eye control vs. brain control) and session type (i.e. static vs. adaptive), we divided each session into training-equivalent trials and test-equivalent trials. We defined training-equivalent trials to be the actual initial eye-control training trials performed at the beginning of static LDA sessions ($n = 379.2 \pm 68.6$ and $n = 322 \pm 6.7$ trials per LDA session for monkeys C and J, respectively) and the first 379 and 322 brain control trials at the start of each adaptive RNN session for monkeys C and J, respectively. Test-equivalent trials were all trials in these sessions that occurred after the training-equivalent trials.

First we looked at saccades during the first 500 ms after GO for correctly decoded training-equivalent trials for both the static LDA and adaptive RNN sessions (Figure 27). Monkey J displayed more erratic saccade behavior during this period, which could either be a behavioral strategy for the task type (i.e. eye- vs. brain-controlled) or the session type (i.e. static LDA vs. adaptive RNN). Monkey C appeared to maintain a similar, direct-saccade strategy during this training-equivalent period for the both session types. We then looked at the first 500 ms after GO for correctly decoded test-equivalent trials (Figure 28). Monkey J's saccades were again more erratic for adaptive RNN sessions than for static LDA sessions, despite the fact that the actual behavior performed after the

GO signal is irrelevant to the decoder. Monkey J's erratic saccade strategy seems therefore more likely due to the session type, and not due to expectations from the session time period (i.e. early vs. late), given the similarity of eye motions between the test- and training-equivalent trials during static and adaptive decoding. Monkey C, again, maintained very consistent, direct saccade behavior and therefore may not have altered its strategy.

While both monkeys displayed individual differences in behavioral strategy in the spatial extent of their saccades, we wondered if either displayed temporal differences. We therefore looked at saccades during the 120 ms express saccade period (Paré and Munoz, 1996), during which there is little possibility that monkeys could react to the trial outcome (Figure 29). Monkey C again appears to have maintained a consistent timing of saccade behavior, and monkey J again appears to have altered its behavioral strategy depending on the trial type. Under brain control, monkey J appears to delays its saccades more than with eye control.

We wondered how erratic saccades were associated with decoder performance, so we compared saccades during correctly- and incorrectly-decoded test-equivalent trials for both static and adaptive sessions (Figure 30). Eye movements during incorrectly decoded trials tended to occur over a wider area and were less stereotyped than during correctly decoded trials. One possible exception to this may be monkey C's static decoding sessions, where we see slightly more erratic activity for correctly decoded trials. In this case, the correctly decoded trials still display relatively stereotyped saccades, however, so it is unclear what conclusions to draw from this.

To verify that erratic saccades were associated with incorrect decoding, we looked at the distribution of saccadic distance from the mean saccades for each of the six targets during the first 0-350 ms after GO for both correct and incorrect test-equivalent trials (Figure 31). Curiously, correctly decoded trials had more divergent saccades than incorrectly decoded trials for both monkeys during the static LDA decoding sessions ($p < 0.001$ for both monkeys C and J; Wilcoxon rank-sum test, Holm-Bonferroni corrected). Notably, this trend reversed for both monkeys during the adaptive RNN decoding task ($p < 0.001$ for both monkeys C and J; Wilcoxon rank-sum test, Holm-Bonferroni corrected) and correctly decoded trials had less divergent saccades than incorrectly decoded trials. To what extent this pattern remained true on a per target basis could potentially elucidate meaning, given the laterality of the implants. We performed a similar analysis, breaking the data down by target (Figure 32) and found that on a per target basis, the relationship between saccade divergence and static decoding performance largely disappeared. Only one target for monkey J continued to show the curious positive relationship between saccade divergence and decoding success during static LDA decoding sessions. On the other hand, we found that the majority of targets continued to show the relationship between saccade stereotypy and decoding performance during adaptive RNN sessions. One ipsilateral and two contralateral targets for monkeys C and J, respectively, showed no significant relationship. In general, correct decoding during adaptive sessions was followed by more stereotyped saccades and there was little relationship between saccade divergence and decoding performance during static sessions. These results help explain why monkey C may have maintained such

consistent saccade behavior during the adaptive decoding sessions, but leaves in doubt the strategy of monkey J, who appears to have had more divergent saccades in adaptive than static decoding sessions. Had monkey J attempted to maintain more consistent behavior during adaptive decoding, it might have enjoyed higher decoding performance.

It is possible that monkey J may not have been attempting a cohesive strategy to solve the adaptive brain control task, potentially due to frustration with the adaptive task paradigm shift eliminating early eye control trials. To determine whether this is a possibility, we looked at the number of fixation errors (both fixation breaks before GO and no fixation before the delay period) committed during the training-equivalent trials for both the static and adaptive decoding sessions (Figure 33). While monkey C's fixation errors did not meaningfully change between the static and adaptive decoding sessions, monkey J's fixation breaks increased dramatically during the training-equivalent trials of adaptive sessions. This false starting could be a sign of either impatience or frustration, or a misunderstanding of the new task paradigm. We wondered if there was a relationship between fixation errors and adaptive decoding performance (Figure 34). We compared both the number of fixation breaks and no-fixation trials to the decoding success rate (which includes only trials in which the decoder made a guess and excludes all fixation break and no-fixation trials) for all adaptive decoding sessions and found that there was only one significant correlation between monkey C's fixation breaks and decoding performance ($\rho = -0.82$, $p = 0.038$, Spearman rank correlation, Holm-Bonferroni corrected). Several other correlations approached significance prior to multiple comparisons correction, including monkey J's fixation breaks and decoding success. The

negative relationship between fixation break rates and decoding success hints that there could either be a drop in decoding performance due to the neural state underlying the increased proclivity to commit fixation breaks (i.e. false starts, possibly) or the fixation breaks are merely a behavioral response to poor decoding. The false starts could indicate impatience or frustration due to the experimental paradigm change of eliminating the eye controlled training trials at the beginning of each session. It is also possible that the false starts were reactions to the poor decoding and did not in-and-of-themselves adversely affect decoder performance.

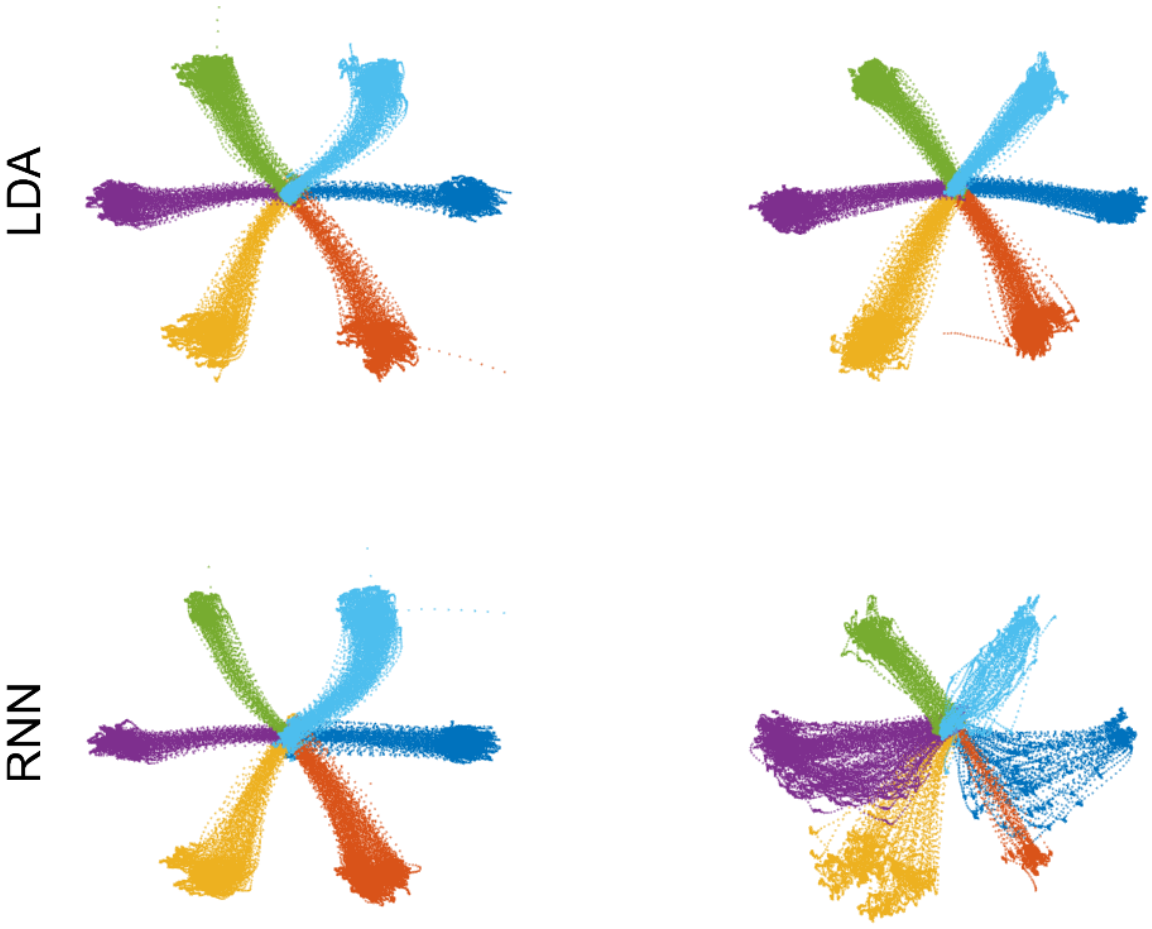

Figure 27: Individual behavioral strategies for eye vs. brain control

The difference between saccades during training-equivalent trials of static and adaptive decoding sessions could indicate behavioral strategy changes due to whether trial is eye- or brain-controlled. Traces represent the nearest 68% of eye motions to the target mean saccades (approximately ±1 STD) for both monkey C (left column) and monkey J (right column) 0-500 ms after GO signal during CORRECT training eye control trials at the beginning of static LDA sessions (top row) and an equivalent number of brain control trials at the start of adaptive RNN sessions (bottom row). Monkey C (left column) maintained similar behavioral strategies to solve both the adaptive decoding task and the eye control task. Monkey J (right column) changed behavioral strategies to solve the adaptive decoding task. Traces are colored by expected target.

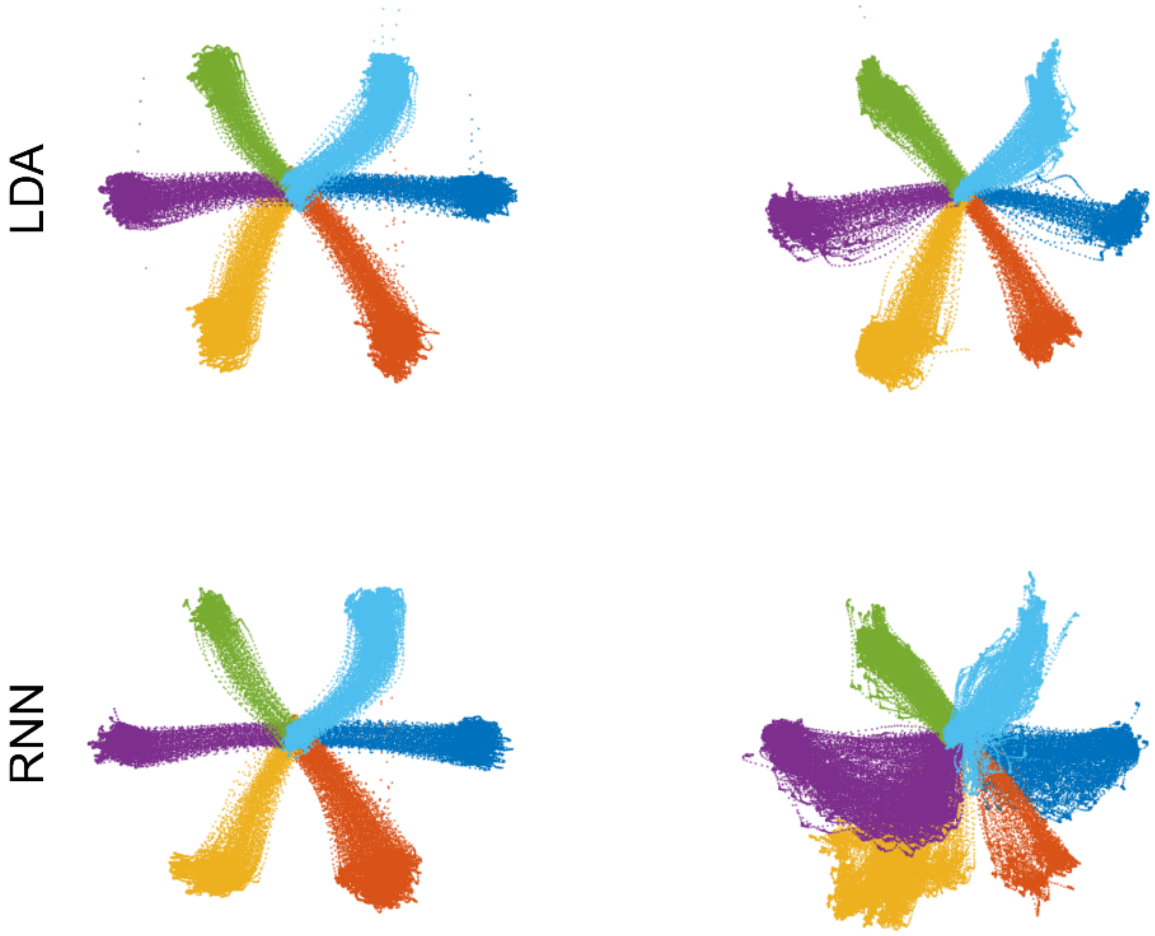

Figure 28: Individual behavioral strategies for static vs. adaptive sessions

Monkeys display individual behavioral strategies between static and adaptive decoders. Traces represent the nearest 68% of eye motions from the target mean saccades (approximately ±1 STD) for both monkey C (left column) and monkey J (right column) 0-500 ms after GO signal during CORRECT test-equivalent brain control trials for both static LDA decoding sessions (top row) and adaptive decoding sessions (bottom row). Monkey C (left column) has very consistent direct-to-target saccade behavior for both the static and adaptive decoding sessions. Monkey J

(right column) has more consistent per-target saccade behavior for the static decoding sessions (top right) and slightly more erratic saccade paths to both ipsilateral (right targets) and contralateral (left targets) targets for adaptive decoding sessions (bottom right). Monkey J's behavioral strategy differences are more likely due to the difference between the two decoding methods, and not due to expectations from the session time period (i.e. early vs. late), given the similarity of eye motions between these test-equivalent trials and the training-equivalent trials during static and adaptive decoding (c.f. Figure 27). It is possible that monkey C's consistent behavior aided in its adaptive RNN performance. Traces are colored by expected target.

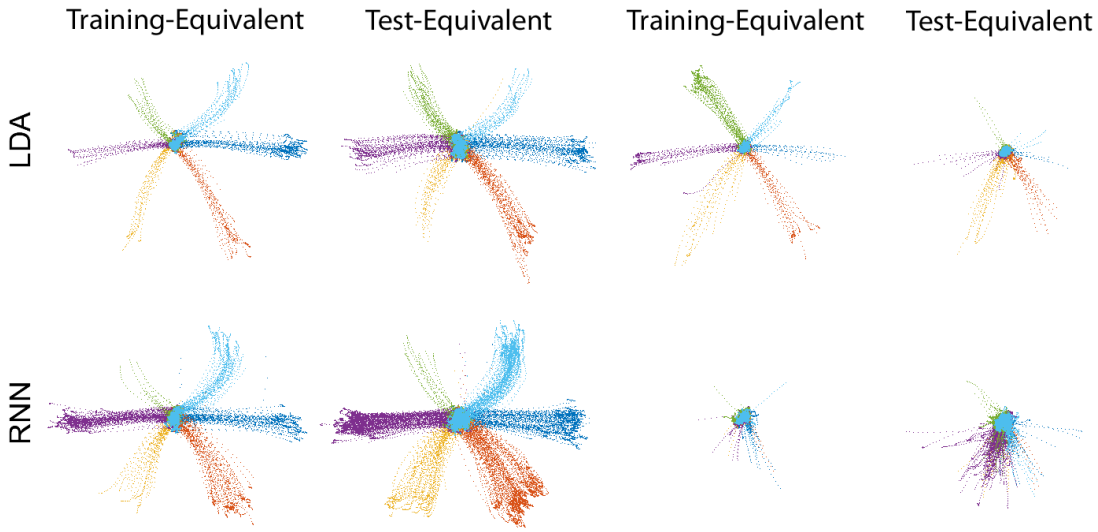

Figure 29: Monkey J changes behavioral strategy under brain control

Traces represent the nearest 90% of eye motions from the target mean saccades for both monkey C (left two columns) and monkey J (right two

columns) 0-120 ms after GO during both CORRECT and INCORRECT training- (1st and 3rd columns) and test-equivalent (2nd and 4th columns) trials for both static LDA decoding sessions (top row) and adaptive RNN decoding sessions (bottom row). Note that monkey C's eye motions 120 ms after GO are qualitatively similar for both static LDA and adaptive RNN training- and test-equivalent trials. Monkey J's eye motions for training-equivalent static LDA trials are faster than its eye motions for any other combination of session or trial types. Traces are colored by expected target.

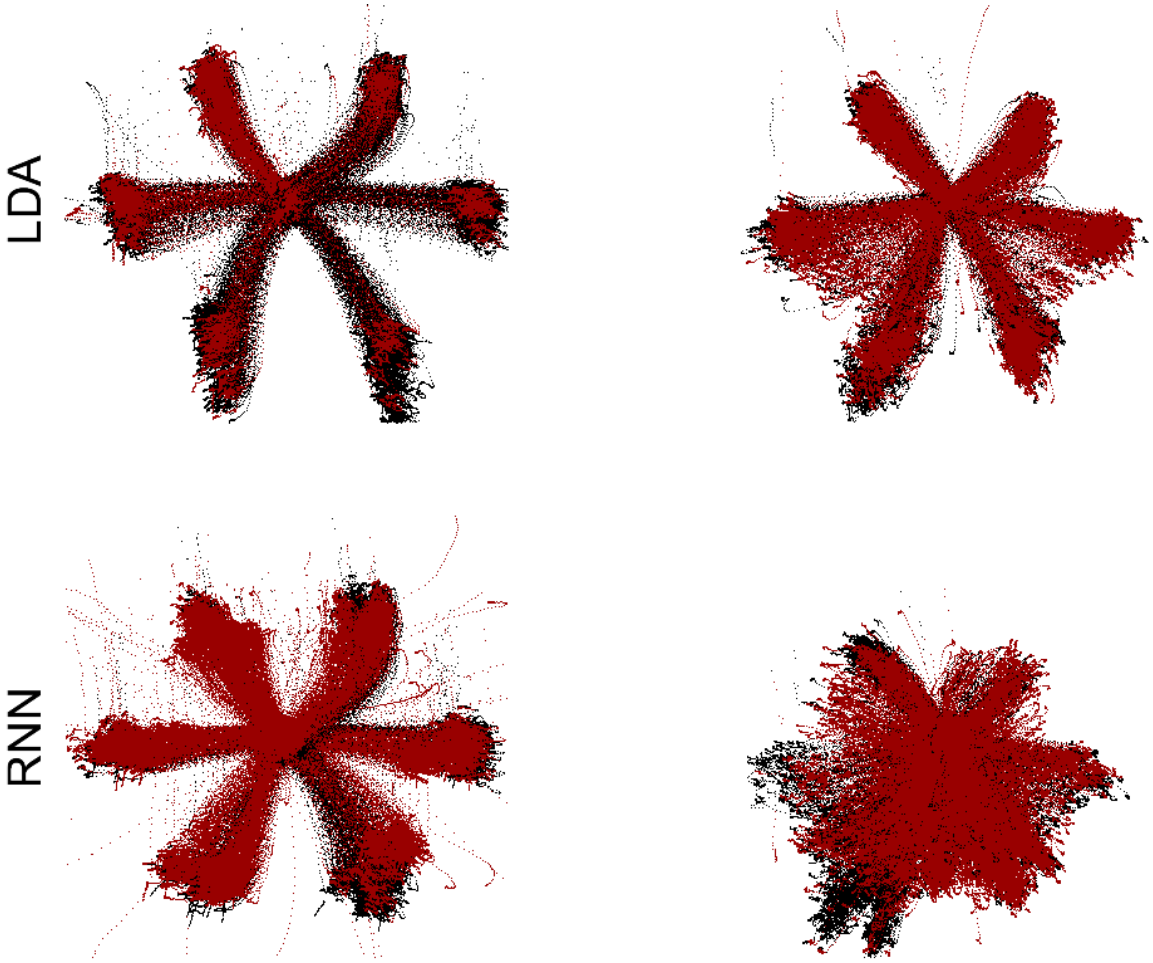

Figure 30: Incorrect decoding responses associated with less stereotyped saccades

Traces represent all eye motions for both monkey C (left column) and monkey J (right column) 0-350 ms after the GO signal during CORRECT (black) and INCORRECT (red) test-equivalent brain control trials for both static LDA decoding sessions (top row) and adaptive RNN decoding sessions (bottom row). Note that eye movements occurring during incorrectly decoded trials tend to occur over a wider area than eye movements during correctly decoded trials. One apparent exception may be monkey C's static decoding sessions (top left).

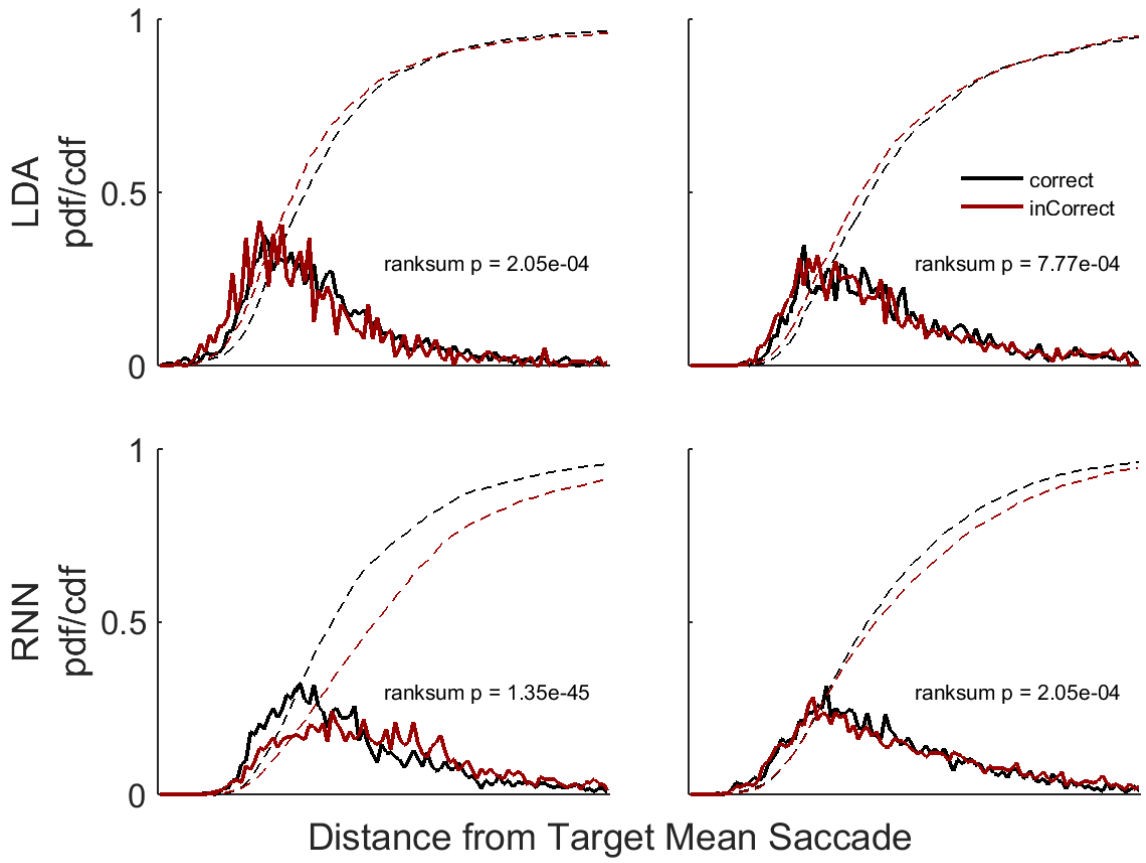

Figure 31: Saccade stereotypy positively associated with correct RNN decoding

We looked at the relationship between saccade deviance from the target mean saccade and decoding performance and found that correctly decoded trials had more divergent saccades than incorrectly decoded trials for both monkeys during the static LDA decoding sessions (top row). Importantly, this trend was reversed for both monkeys during the adaptive RNN decoding task (bottom row) and correctly decoded trials had less divergent saccades than incorrectly decoded trials. All axes in all plots are equal. All comparisons were made using the Wilcoxon rank sum test, and p-values

were Holm-Bonferroni corrected. Note the pdfs are scaled for clear viewing.

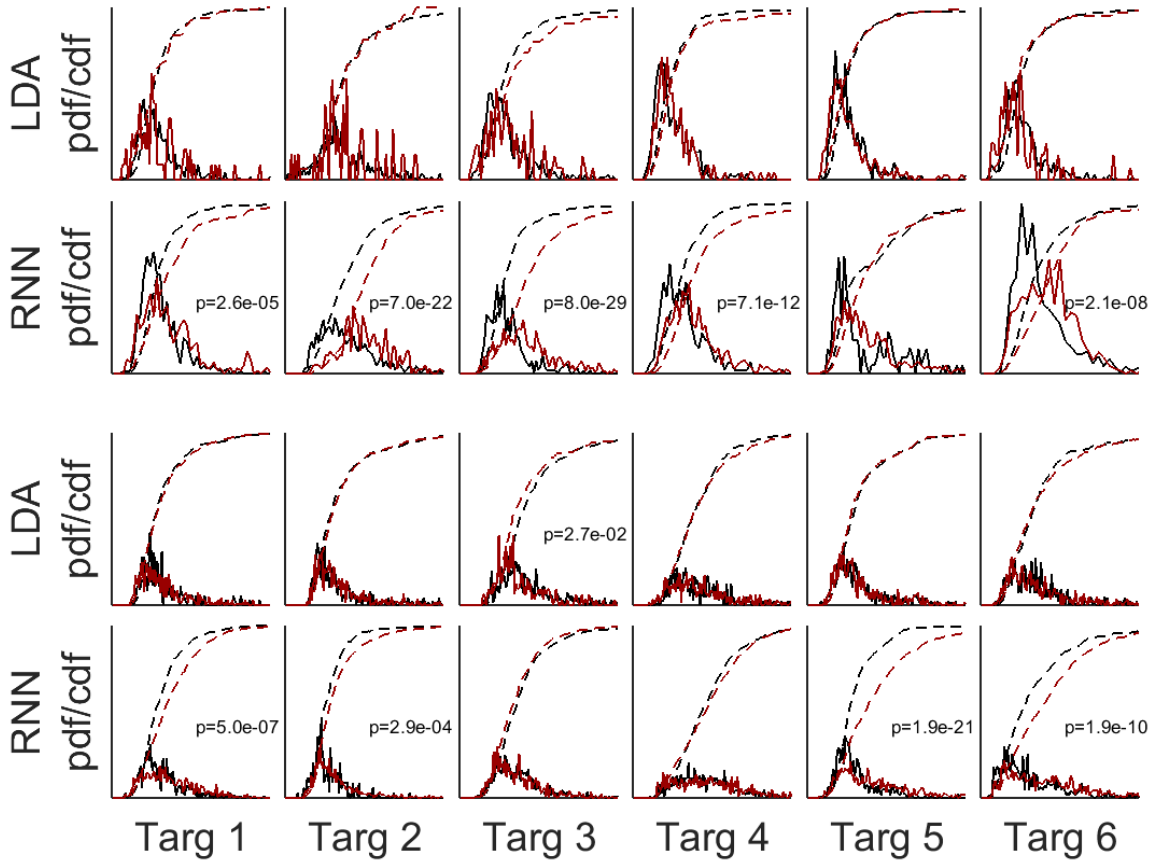

Figure 32: Positive performance-stereotypy association specific to RNN decoding

Distance from target mean saccade has larger detrimental effect on adaptive decoding than static decoding. We performed similar analysis as in Figure 31, broken down by target to determine whether there is a target-specific effect, given the laterality of the implants. In general, static decoding sessions (1st and 3rd row) did not have significant differences between correct and incorrect saccade deviations on a target-specific level,

but adaptive decoding sessions (2nd and 4th rows) did. Top two rows: monkey C. Bottom two rows: monkey J. All comparisons were made using the Wilcoxon rank sum test, and p-values were Holm-Bonferroni corrected. Note the pdfs are scaled for clear viewing.

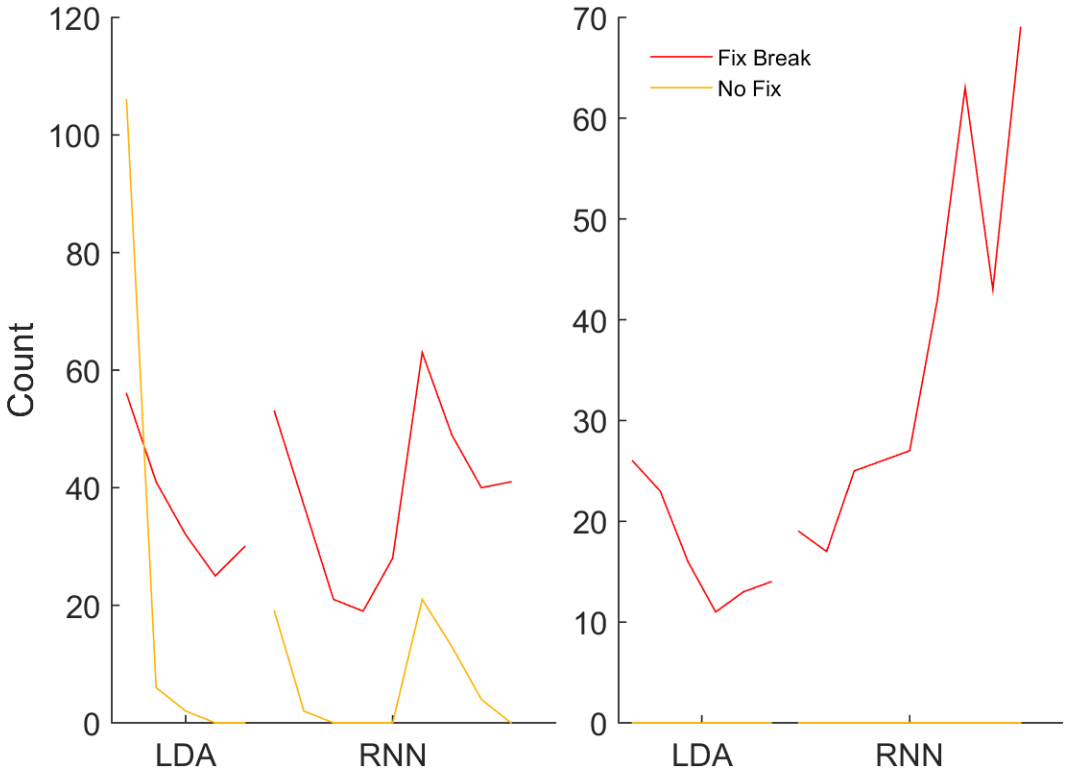

Figure 33: Monkey J has more fixation errors during adaptive sessions

Monkey J (right) increases the number of fixation break errors during the training-equivalent trials (~350) of adaptive sessions. This false starting could be a sign of impatience or frustration. Monkey C (left) did not meaningfully alter its fixation error trends between static and adaptive

decoding sessions. Note the oscillations in the data for monkey C correspond to the weekly session cycle.

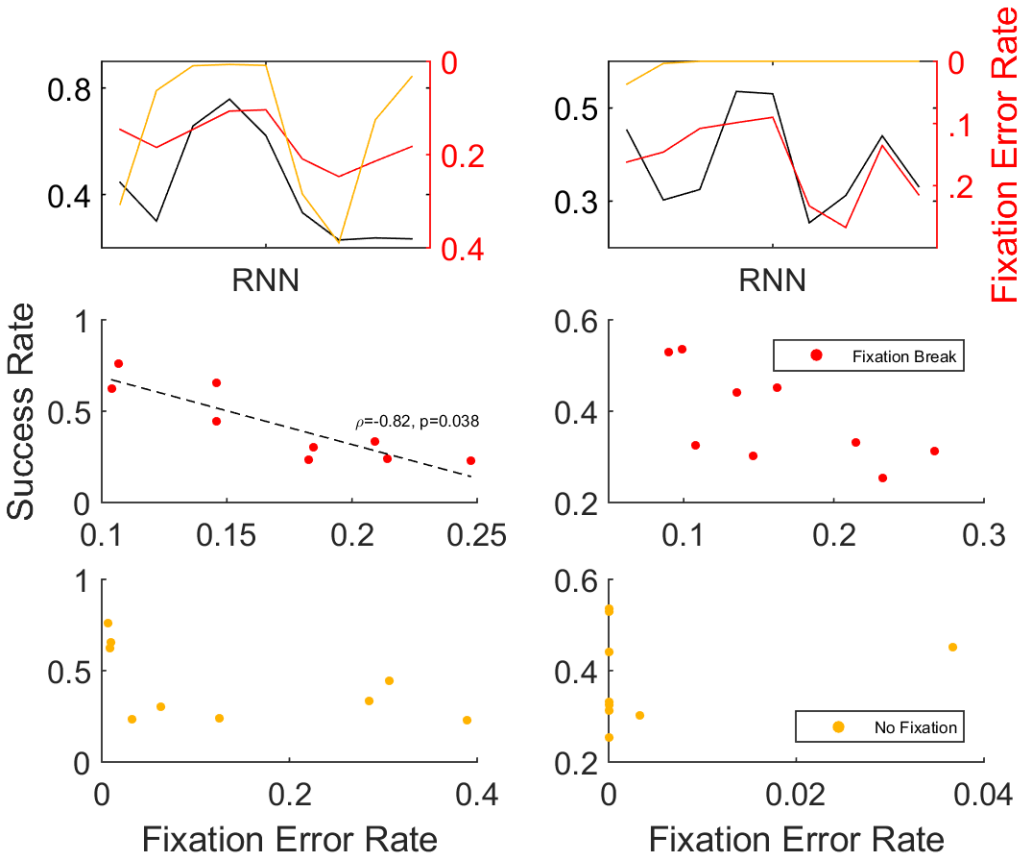

Figure 34: Fixation break rate inversely correlated with decoding performance

Top row: Success rate (left y axis) and fixation errors (right y axis; note the values decrease) for Monkeys C and J (left and right). There is an apparent negative correlation between the fixation errors and decoding performance. Left middle: a significant negative correlation was found for the fixation break rate and decoding performance for Monkey C ($\boldsymbol{\rho = -0.82, p = 0.038}$; Spearman rank correlation, Holm-Bonferroni corrected). Right middle: a negative correlation was found for the fixation

break rate and decoding performance for Monkey J, but did not survive multi-comparison correction. Bottom row: No significant correlations were found between no fixation rate and decoding performance. These results suggest that monkey frustration was exhibited through fixation breaks.

Reconstruction of 80-475 Hz Power

In the learning study, we studied the feasibility of using beta band power as a "switch" for brain machine interfaces. One of the reasons for selecting this band was the robustness of low frequency neural measurements, in the face of mechanical or biological changes in the recording system. High frequency activity above ~70-80 Hz is purported to be a marker of non-linear neuronal activity, such as spiking (Manning et al., 2009; Miller et al., 2009a), and could be a more robust decoding feature than spikes, so it was reasonable to use this neural feature for decoding. In contrast, lower frequency power in dlPFC, SEF, and FEF was found not be as informative for decoding saccades in related work (Brincat et al., 2013a).

A potential reason for a lack of decoding success using low frequency LFPs is that they may be representative of top-down signaling (Engel and Fries, 2010; Buschman et al., 2012; Friston et al., 2015) or are otherwise associated with "input" to an area, rather than "output" from an area, which may be more closely tied to high frequency LFPs or MUA (Logothetis, 2003; Gao, 2016). We wondered if we could leverage the robustness of low frequency LFP measurements while increasing the information extracted from them by modeling cross-frequency interactions between low and high

frequencies and reconstructing the high frequency "output" activity (Brincat et al., 2013a). This way we may be able to "encode" some of the underlying transformations between low frequency input to an area and the 80-475 Hz marker of output from the area, which was shown to be the most informative frequency band of saccade direction during the delay period (Brincat et al., 2013a).

We performed phase-amplitude (**PAC**) and amplitude-amplitude (**AAC**) cross-frequency coupling analyses of several low frequency bands to an offline measurement of 80-475 Hz power. For all sessions of each monkey, we selected random pairs of electrodes and calculated phase-amplitude coupling in 1 Hz bands from 1 – 5 Hz to 80-475 Hz power, and amplitude-amplitude coupling of the 1 – 2, 2 – 4, 4 – 7, 7 – 12, 12 – 20, 20 – 30, and 30 – 50 Hz frequency bands to 80-475 Hz power. These CFC analyses are not rigorous and were performed simply to determine potential decoding features for recreating 80-475 Hz activity; detailed CFC analyses are left to colleagues in ongoing work (Jia, dissertation). Monkey C commonly displayed 2-3, 3-4, and 4-5 Hz PAC and 1-2, 12-20, and 30-50 Hz AAC, and monkey J commonly displayed 1-2 and 2-3 Hz PAC and 1-2, 2-4, and 12-20 Hz AAC. We calculated the phases and amplitudes of these PAC and AAC bands for each electrode during the ten 75 ms time steps of the delay period for 80-475 Hz encoding. Though studies have shown the utility of phase in determining motor intent (Rubino et al., 2006), several iterations of development led us to eliminate PAC inputs and further narrow our feature selection to the 1 – 2, 12 – 20, and 30 – 50 Hz AAC bands for monkey C, and the 1 – 2, 2 – 4, and 12 – 20 Hz AAC bands for monkey J.

We used a simple feedforward artificial neural network (**ANN**) with one hidden layer (100 logistic units) to reconstruct 80-475 Hz activity for each of the 96 electrodes during the delay period. We calculated the RMS error (**RMSE**) and found both the validation and test set results closely matched, indicating a lack of overfitting (Figure 35). How the RMSE values translate to reconstruction accuracy depends on overall 80-475 Hz power of the electrodes. We compared actual and reconstructed values during a high RMSE session (monkey C's 2nd adaptive RNN session RMSE ~4) and a low RMSE session (monkey J's 2nd adaptive RNN session RMSE ~1) and found that although temporal fluctuations within the delay period are lost, overall 80-475 Hz power reconstruction for a given trial was excellent for all electrodes, despite the apparent range of RMSE values (Figure 36). Versions of the ANN with recurrent connections did not improve the fidelity of the temporal fluctuations.

The 1-2 Hz band appears to be especially important for reconstructing 80-475 Hz power, based on the largest 0.5% of network weights (by magnitude and excluding the bias weights) in each layer of the ANNs (Figure 37). The inputs for this band consistently had many of the largest outgoing weights to the hidden layer across both monkeys, many sessions, and many of the 20 generated reconstruction models for each session. This could suggest that inputs to dlPFC, SEF, and FEF oscillating between 1-2 Hz are consistent, large factors of 80-475 Hz output activity in these cortical areas. Another common finding is that within each session there were several consistent output units with large incoming weights for many of the reconstruction models generated for that particular session, which could be because of those electrodes' high 80-475 Hz power.

An example of this can be seen in Figure 37, where monkey C's electrode 35 (in SEF) with the extreme power has many of the largest incoming weights from the hidden layer.

This reconstruction work was intended to increase the information available to RNN decoders. We hoped to accomplish this in one of two ways: (1) use the hidden layer activations as an "encoder" of the low-frequency-input to high-frequency-output transformation to either be used as the inputs or as additional inputs to RNN decoders as previously described, or (2) use a reconstruction ANN and an adaptive RNN as the first and second stages of a deep neural network.

Preliminary attempts were made to address both of these possible methods, with limited success. First, we used the reconstruction ANN's hidden layer "encoder" as both inputs and additional inputs to the previously described RNN architecture, with little decoding performance. Second, we made partial attempts to test a deep neural network and reasoned that if we were able to reconstruct the 80-475 Hz power using low frequency power with reasonably high fidelity, then it should be possible for a RNN with the equivalent number of layers to learn the weights for the additional step of decoding the reconstructions to guess the saccade target. This was also not successful, but we did not properly pre-train this deep neural network as described in (2). The lack of success in using low frequency activity to decode saccades could either be due to a genuine lack of information in these frequency bands (Brincat et al., 2013a) or to inappropriate RNN architectures or learning parameters. Either method could potentially provide more decoding power but, due to time constraints, will have to be investigated in future work.

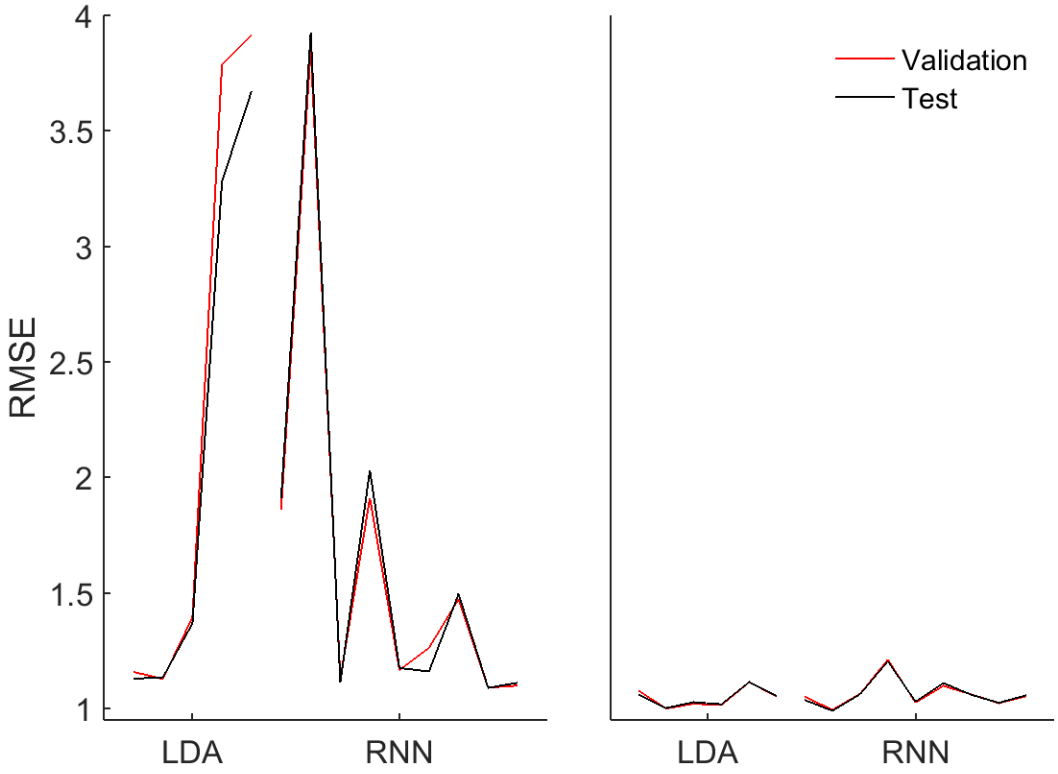

Figure 35: Reconstruction models did not overfit data

Close validation and test RMSE values indicate that the reconstruction models did not overfit the data.

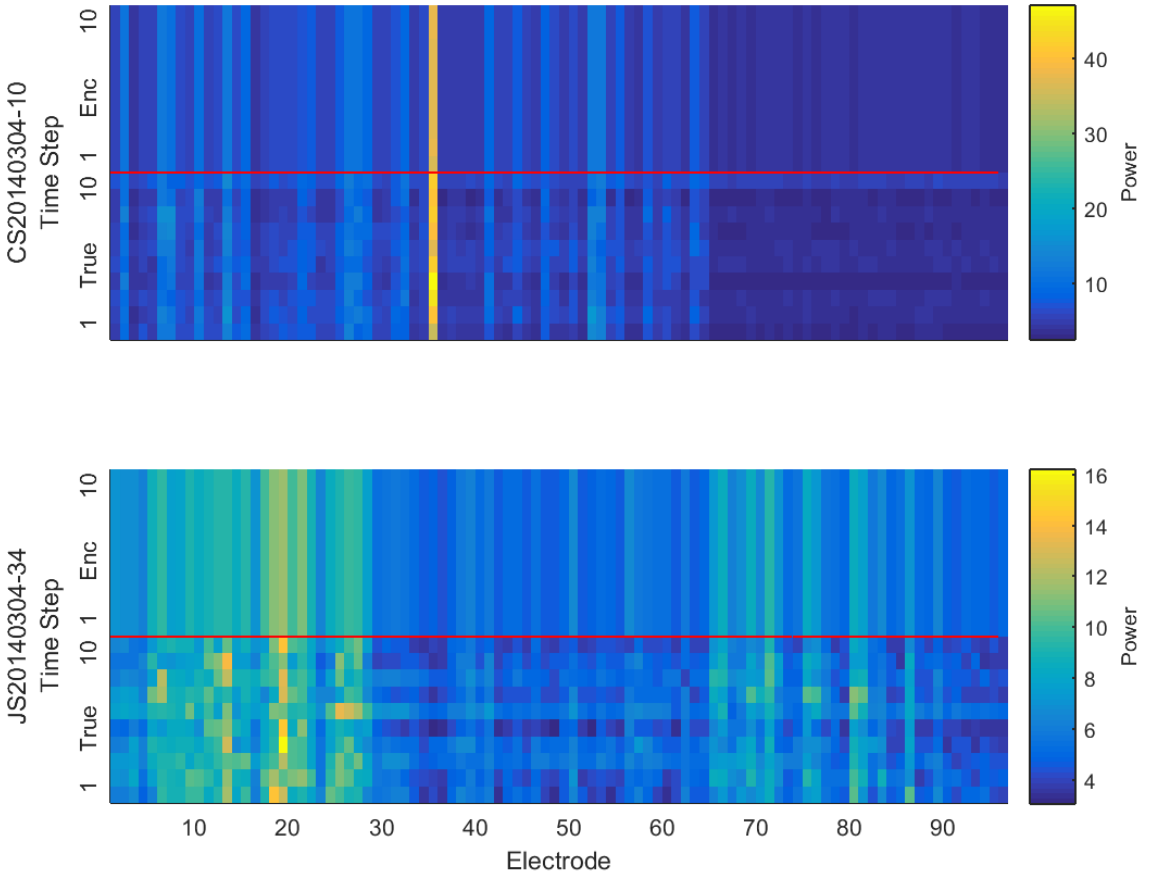

Figure 36: Good 80-475 Hz reconstruction accuracy

Example of 80-475 Hz electrode power reconstruction demonstrates excellent accuracy of encoding model and that the RMSE measures are likely affected by absolute power levels in the neural recordings. Note electrode 35 (in SEF) for monkey C (top) has relatively high power, but that overall reconstruction of the electrodes is good. The overall power of the electrodes for monkey J (bottom) is lower, and reconstruction is still good. The temporal variation is not encoded with high fidelity, but overall electrode power levels are captured accurately. Top: 80-475 Hz power reconstruction for trial 10 from monkey C's session 20140304. Bottom:

80-475 Hz power reconstruction for trial 34 from monkey J's session 20140304.

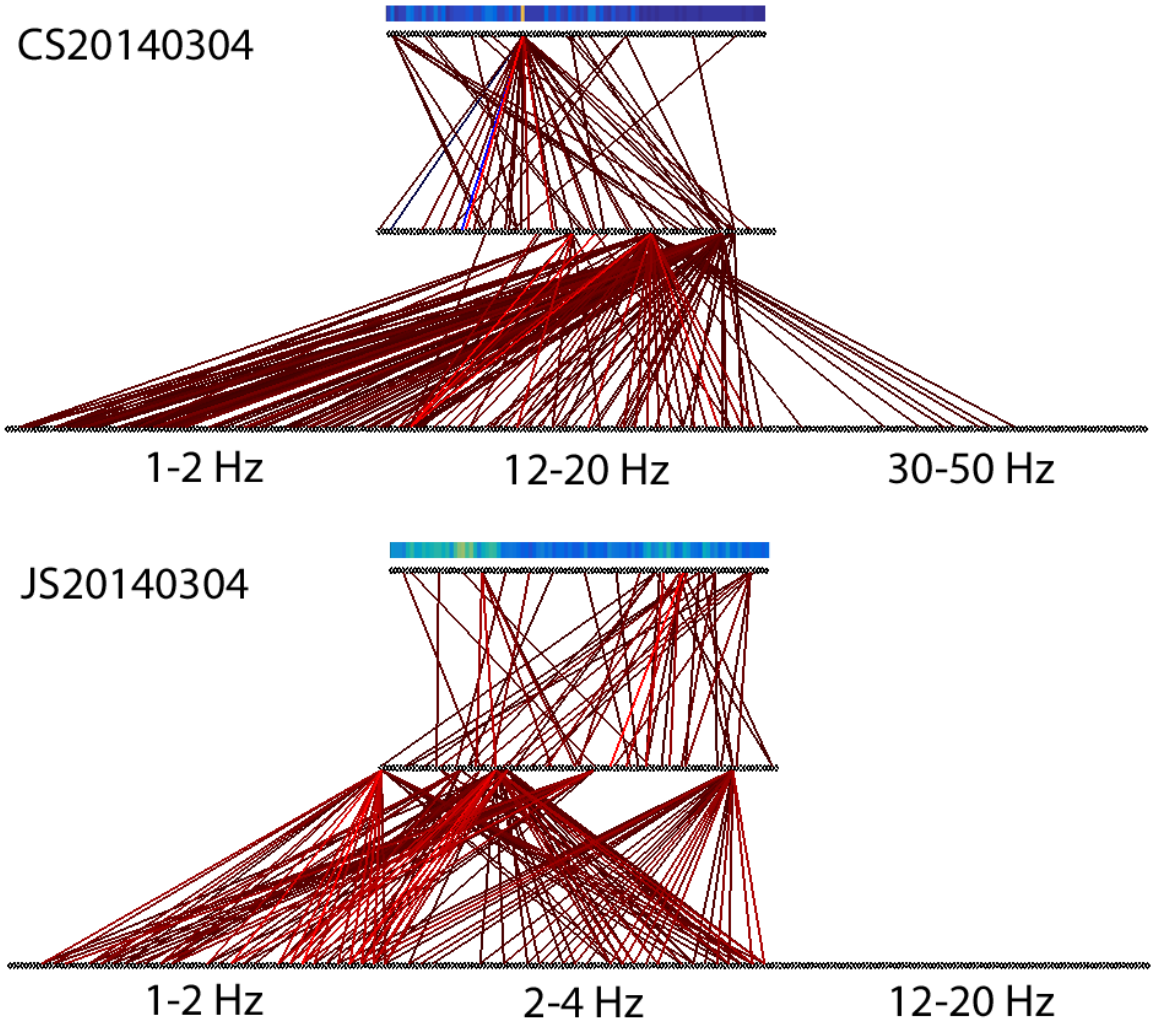

Figure 37: 1 – 2 Hz LFP power consistently a large factor in 80-475 Hz power

Largest 0.5% of network weights (by magnitude) suggest 1-2 Hz power is a consistent, large factor of 80-475 Hz power. Shown above are each 1 of 20 encoding models generated for each of the same two sessions shown in Fig. ENC EXAMPLE. Top: largest 0.5% of weights (excluding the bias

weights) in each layer for 1 of 20 encoding models from monkey C's session 20140304. Bottom: largest 0.5% of weights (excluding the bias weights) in each layer for 1 of 20 encoding models from monkey J's session 20140304. A common finding from many of the encoding models across many of the sessions is that the 1-2 Hz power inputs have many of the largest outgoing weights. Another common finding is that there are several output units from each session with similarly large incoming weights for many of the encoding models generated for that particular session. At the top of each plot is one reconstructed time step from the same examples shown in Fig. ENC EXAMPLE. Note that for monkey C (top), electrode 35 with the extreme power (in SEF) has many of the largest incoming weights from the hidden layer, suggesting that the commonality of several output units frequently having the largest incoming weights could be due to high 80-475 Hz power levels.

Discussion

We showed the feasibility of a generalized simple RNN to adapt to novel trial data offline. Closed loop adaptive RNN decoding performance was considerably lower than both offline RNN and closed loop static LDA decoding performance. Contralateral targets were decoded more successfully than ipsilateral targets in both offline and closed loop RNN decoding, which is likely a direct consequence of the unilateral implants (Bruce and Goldberg, 1985).

Closed loop calibration trials could be eliminated through the creation of pre-trained biomimetic decoding models, which both our offline and closed loop results support. Initial performance of both the offline and closed loop decoding models was above chance in most cases. While offline adaptive performance consistently increased throughout a session, closed loop adaptive performance sometimes increased but most often remained flat or decreased.

While the data do not specifically address whether the adaptive RNN would adapt to fundamental changes in the neural measurements, the data suggest that it could. Initial performance of the offline decoding models was worse than previous day batch performance of those same models, indicating that changes occurred in the neural measurements between days (e.g. electrode drift or neural plasticity). If a change had not occurred at some point between the previous and next day's trials, then those offline models would have performed similarly. Performance thereafter increased, suggesting that a closed loop model is likely to adapt to whatever changes occurred—in the absence of subject awareness to the decoding task and model.

We should not be surprised that our closed loop results were different than offline results and that the subject-decoder interaction was detrimental to the closed loop adaptation process (Suminski et al., 2010). Given the biomimetic nature of the decoder and the good performance of the offline simulation, we suspect that poor closed loop performance was caused by subject frustration or confusion with the elimination of preliminary eye control training trials (Orsborn et al., 2011). Prior to the closed loop

RNN sessions, eye control determined task success and reward for 300 successful trials and only afterwards was brain control required. For the closed loop RNN sessions, brain control started on the first trial. This alteration to the task paradigm may have frustrated the monkeys, decreasing motivation and task participation. Subject motivation has been shown to be vital to decoder performance (Musallam et al., 2004). Initial closed loop performance was comparable to initial offline performance, which suggests that any differences between offline and closed loop adaptation performance are more likely to be due to subjects' responses to the decoder than to improper training of the initial batch models.

Whether the recorded changes seen in the neural data were of biological or electromechanical origin is unknown, but the correlation of 80-500 Hz power between days sometimes changed significantly. Given the relative inter-session stability of 80-500 Hz electrode power correlations between the static LDA sessions and the higher variability of these correlations starting with the adaptive RNN sessions, it is possible that subjects altered their neural control strategies in response to the adaptive decoding task itself or to the task paradigm change. That Monkey C's FEF activity had the largest change after the first week of decoding possibly suggests the possibility of a neural strategy change of saccade "effort", as FEF disruption was found to adversely impact saccade magnitude (Cameron et al., 2015).

It is possible that the learning parameters of the decoder were not appropriate, considering the ability of the subjects to learn (Ganguly and Carmena, 2009). Even though we generated successful decoding models using identical learning parameters and

RNN architectures using data sets 104 days apart in the preliminary offline analysis, the potential for uncontrolled co-adaptation exists if the RNN weights changed too rapidly. Either the subjects or the decoders may have been overcompensating with their respective correction mechanisms, given the difference between effective learning of static decoders (Ganguly and Carmena, 2009) and effective learning during closed loop decoding (Kim et al., 2008; Chase et al., 2009; Koyama et al., 2010; Cunningham et al., 2011).

Anecdotal reports suggest that simpler decoding algorithms with more direct causes and effects, such as population decoding, are easier to learn by BMI subjects. Linear decoders are capable of adapting to non-stationary neural signals (Gürel and Mehring, 2012), but it is unclear if simpler decoding methods will result in better long-term control and recurrence could allow better decoding of motor sequences (Jordan, 1986). Until we arrive at a comprehensive theory on motor control, we assume the possibility of efficient control through complex decoding methods, given enough time to learn. This is analogous to healthy humans who generate movements despite having no conscious control over specific neurons and do not achieve skilled motor control for years.

The two monkeys displayed opposing behavioral strategies. Monkey C's saccades after GO were both spatially and temporally consistent, whereas monkey J's saccades were more spatially inconsistent for adaptive RNN sessions and quicker during eye control trials. Monkey C's strategy may therefore have been to simply maintain consistent delay period neural activity, no matter the trial conditions. On the other hand,

monkey J may have adjusted strategies for both eye- vs. brain-control trials and static vs. adaptive sessions. Monkey J's saccades were more spatially erratic for the adaptive sessions, which could either be a sign of attempted learning of the decoder or frustration due to the elimination of eye control trials. Additionally, monkey J's saccades were faster during eye control, many of which occurred within the express saccade period (Paré and Munoz, 1996). Despite that the behavioral response after GO should have no direct bearing on the decoder response from the delay period activity, the delay period activity should, on the contrary, have some bearing on the behavioral response. Therefore, why monkey J appears to have segregated its behavioral response profile into spatial (static-decoder-direct-saccade / adaptive-decoder-erratic-saccade) and temporal (eye-control-fast-saccade / brain-control-slow-saccade) modes is unclear.

Addressing the spatial difference in saccades, the mathematics of the adaptive RNN and the static LDA decoders is vastly different. The RNN could then very well have learned different features or applied different meanings to those features than what the LDA decoder did. Once monkey J switched to the RNN decoder after being accustomed to the LDA decoder, the apparently erratic saccades may actually have been the outcome of strategic attempts to find appropriate delay period neural activity that would better align with RNN feature preferences. Temporally, it is possible that monkey J learned to better control the static LDA decoder using slightly delayed saccades and continued this learned behavior with the adaptive decoder, rightfully not respecting or understanding the difference in the underlying decoding methods. All in all, although monkey J displays more behavioral differences that could have been learning attempts, monkey C's

consistent behavior may have been partially responsible for its higher adaptive RNN performance than monkey J.

The timescale of feedback is a factor in subject learning. Few adaptation studies utilized discrete classification tasks in which feedback is only provided at the end of a trial (Bryan et al., 2013; Bishop et al., 2014). In our case, the monkeys only received feedback at the end of a 750 ms delay period, which reduces the association between neural activity and desired outcome. Normal visuomotor and sensorimotor feedback processing times, in contrast, occur on the order of ~100ms (Zelaznik et al., 1983; Liu et al., 2005; Pruszynski et al., 2011). One could imagine that monkey J's delayed saccades after GO during brain control was a strategy to continue delay period activity through feedback to better connect the underlying brain activity to the decoding outcome. It may also be that monkey J's motor activity was reduced during brain control, as subjects in previous BMI studies have also done (Taylor et al., 2002; Velliste et al., 2008). Why monkey C did not appear to explore the behavioral and neural feature space as readily as monkey J to solve the brain control task is unknown. Monkey C had overall higher performance, however, and we showed that stable neural activity was advantageous, so it may have been that monkey C was attempting the best strategy.

The rate and timescale of adaptation is also a large factor in decoder performance. Though Orsborn et al. conclude that decoder adaptation is best applied in mini batches on an intermediate time scale of several minutes to allow the subject to co-adapt, there is no obvious reason why continuous, slow learning could not also be similarly advantageous

(Wu and Hatsopoulos, 2008; Orsborn et al., 2011, 2012). The learning rate of an RNN could be set low enough that a subject should be able to co-adapt readily. After all, it is likely sudden, drastic changes in decoding outcomes that defy expectations which cause subjects to alter neural patterns and realign their forward models. This again could help explain why the session paradigm change of eye control trial removal could have been disruptive enough for the monkeys to overcorrect.

The lack of success utilizing the high frequency reconstruction ANN hidden layer "encoder" as inputs to the previously described RNN architecture is surprising, given the high fidelity of the reconstructions. It is possible that the inability to recreate the small temporal fluctuations in power is a sign of a much larger problem than it would seem. The "encoder" that appears to perform so well may, in fact, not be modeling the underlying neural processes converting low frequency input to high frequency output with any accuracy whatsoever. Indeed, the bias weights are almost uniformly the largest weights going from the hidden layer to the output layer, suggesting a common baseline of 80-475 Hz power, with only minor variation based on the target. Visualization of the per-target electrode activations partially confirms this. Delta band EEG power was found to positively correlate with MUA activity (Whittingstall and Logothetis, 2009), so it is reasonable that the next most common group of large weights were from the 1-2 Hz amplitude inputs to the hidden layer. There is only ~1 cycle of 1-2 Hz during the 750 ms delay period window, which means there is not likely to be large variations in 1-2 Hz amplitude during this time. Had we kept 1-2 Hz phase as an input, we might have better

reconstructed the variations of 80-475 Hz power within the delay period (Whittingstall and Logothetis, 2009).

It is also possible that different RNN learning parameters need to be used due to the different methods of closed loop and offline power calculation and the range of their values (the closed loop method generally produced lower values). However, given the relative accuracy of the reconstructions, the two-part deep neural network may be a possibility and could help reduce training time of such a network. Additionally, deep neural networks may better capture natural dynamics that could lead to improved decoding performance (Guclu and van Gerven, 2015; Sussillo et al., 2015).

CHAPTER FOUR: Integrating Learning and Adaptation

Results of Learning and Adaptation Studies

The learning study had mice controlling an auditory cursor, by modulating beta band power in vibrissal motor cortex (Torene et al., 2015). Learning to control beta band power using auditory neurofeedback would be a more effective feedback modality than visual feedback for visually impaired people, including some ALS patients, assuming low dimensional control (McCreadie et al., 2013). We studied the neural and behavioral correlates of learning in an LFP-based decoder in the context of motor state detection (Mason and Birch, 2000; Bai et al., 2008; Fatourechi et al., 2008; Wang et al., 2012). Previous studies failed to distinguish between motor imagery and motor execution (Birch et al., 2002), which is problematic because LFP dynamics underlying imagined movement look very much like the dynamics underlying actual motor movement (do Nascimento et al., 2006; DaSalla et al., 2009; Miller et al., 2010). If we are able to train subjects to volitionally control beta SMR power, it could serve as an appropriately high SNR motor intent signal distinguishable from motor imagery. This would allow an inflated threshold for BMI execution, which could reduce the number of false positives occurring during motor imagery or rest. Furthermore, while several studies have made behavioral observations of subjects during BMI learning (Lebedev et al., 2005; Velliste et

al., 2008; Moritz and Fetz, 2011), few have thoroughly analyzed the behavioral correlates of learning to modulate LFPs.

Mice in our neuromodulation learning study increased performance quickly over the course of a few sessions. They were able to perform the task via production of beta range burst activity, and not through behavioral strategies. The beta range bursts increased in duration and not in amplitude or rate of occurrence, although these last two features likely contributed to performance. MUA in vM1 appeared to drive high frequency EMG activity in the beta range during task-related periods.

The adaptation study applied adaptive decoding methods to predict saccade direction in monkeys (Brincat et al., 2013b; Torene et al., 2013). Only very recent work has attempted to decode similar activity (Graf and Andersen, 2014; Ohmae et al., 2015; Boulay et al., 2016). The focus of the current study is the adaptive component for discrete classification in a closed loop setting. A previous discrete offline classifier was adaptively trained on a daily basis to predict reach direction, which showed that using decoding parameters from a previous day could help or hinder initial performance on the next day (Bishop et al., 2014). Previous studies have used adaptive artificial neural networks in co-adaptive paradigms, although either generalization was problematic, or the decoding was continuous rather than discrete (DiGiovanna et al., 2009; Sussillo et al., 2012). We investigated the use of simple RNNs in a discrete decoding paradigm, and we implemented several generalization techniques to provide the simplest, most biomimetic decoder on which to make small adaptations to maintain consistent performance (Danziger et al., 2009; Sadtler et al., 2014). After promising results from offline

simulation, we tested our classifier online with two macaque monkeys. The classifier was recalibrated after each trial, but after brief improvement performance became poor. The causes of the poor performance are unclear, but are vital to our understanding of how to solve the "two-learner" problem (DiGiovanna et al., 2009; Shenoy and Carmena, 2014), in the case of co-adaptive BMI systems (Dangi et al., 2013; Merel et al., 2015).

Whereas offline simulation of the adaptive RNN decoding was successful, closed loop adaptive RNN decoding was less successful. As may be expected, we found that maintaining similar neural activations and behavioral responses was associated with higher decoding performance. The monkeys' awareness of closed loop decoding likely altered their learning strategies, especially in light of the session paradigm change in which initial eye control training trials were eliminated from the adaptive decoding sessions.

We studied neural and behavioral responses of the mice and the monkeys in the learning and adaptation studies and found that the two cohorts had varying levels of control strategy similarity. While the two monkeys displayed several neural and behavioral strategies in a 6-choice delayed saccade task, mice had highly similar behavioral correlates of successful task performance. Given broad freedom to solve the task using whichever strategy of their choosing, all three mice performed successful trials roughly as follows: (1) poke the snout into the trial initiation port, and (2) remain within the trial initiation port while continuously performing various orofacial behaviors, including whisking, licking, biting, and sniffing, until the reward tone sounded. Throughout this basic procedure, the vM1-EMG relationship—as quantified by task-

related LFP-EMG coherence changes and MUA-EMG spike-field coherence—appeared highly similar among the three mice. No "action" was being decoded—the mice had no requirement other than to increase the BPR value, yet the vM1-EMG relationship appeared relatively consistent between the three mice. In contrast, monkeys engaged in various neural and behavioral strategies, even though it was shown that consistency in both was associated with higher decoding success.

In the case of the mice, one possibility may explain the apparent paradox: the requirement to increase the BPR value may not be as generic as we thought. Beta band activity is associated with neural states and is a more global signal than higher frequency activity, but there may be more than one beta rhythm—not all of them volitionally controllable. It is possible that beta band activity that is possible to neuromodulate may be a more "intentional" beta that occurs within a narrow behavioral context. For example, preparatory-period-motor-command-adjustment-beta (Torrecillos et al., 2015) and tonic-hold-beta (Baker et al., 1997) may be more intentional activity or is otherwise more "controllable", whereas error-salience/forward-model-updating rebound-beta (Torrecillos et al., 2015; Cao and Hu, 2016; Tan et al., 2016) may be a more automatic process underlying neural plasticity and motor learning. The consistent vM1-EMG relationship and BPR increases during high EMG suggests the likelihood that the modulated beta rhythm was the tonic-hold-beta of the former beta description. This directly implicates tonic hold of those muscles for which the vM1 cortical area controls, implying the inter-mouse behavioral similarity.

Neural and behavioral consistency from the monkeys would likely have improved decoder performance. The decoder would not have committed any errors if the monkeys' neural activity was identical within each of the 6 targets, but this is of course not possible for many reasons (Faisal et al., 2008). Additionally, the electrode arrays were implanted in oculomotor and working memory areas, but because the decoding signal is high frequency activity that is likely to be of highly local origin, the exact placement of each electrode on the arrays may not be comparatively good between the monkeys for the specific targets. Monkey J had lower performance than monkey C throughout many earlier decoding sessions prior to this adaptation study, so exploratory behavior may have been the best strategy based on a circumstance of the electrodes—while consistency would have been advantageous during the course of the data collection period of the adaptation study, more effective learning may be achieved over the long term through exploratory neural "noise" (Krogh and Hertz, 1992b).

The timescale of feedback and the size of neural integration windows is an important factor in how well subjects are able to learn BMI control (Cunningham et al., 2011). Cunningham et al. (2011) found that whereas offline decoding simulations worked best using 100-200 ms neural integration windows, closed loop decoding worked best with 25-50 ms windows. This difference is very likely due to whether subjects' continuous alterations of control strategy in response to decoder feedback occur within the neural integration window before the decoder can update its estimation of the neural state. Shorter windows increase system noise, for which subjects can compensate in a closed loop scenario; longer windows reduce the association between intent and decoder

estimate feedback, to which a subject is oblivious in an offline scenario. One recent offline saccade decoding study found that using the average firing rates of neurons across all task epochs, including the cue, delay, and response period had the highest decoding accuracy (Boulay et al., 2016). The results of this study are confounded by many of the problems outlined above, namely: (1) an offline study that does not account for subject interaction, (2) utilizing extremely large neural integration windows, and (3) using spiking activity, which may not be a valid long term solution. In our adaptation study, intended saccades were decoded using a 750 ms neural integration window (i.e. the 750 ms delay period) and only provided one discrete target estimate as feedback. Were the task of a more continuous nature (e.g. a saccade sequence; Shanechi et al., 2012), the monkeys would have received more feedback in less time and might, therefore, have had different control strategies. As it was, monkeys could only guess in hindsight what it was they did that was more "correct" or "incorrect".

In contrast to discrete target selection, the mice in our neuromodulation learning study received continuous feedback, albeit over a 600-1000 ms neural integration window. Recognition of LFP oscillations, in the case of the BPR decoding feature, by its nature requires more than one cycle of the oscillation to occur, so the longer window may not have had much negative effect in this particular task. Further, longer neural integration windows with shorter feedback windows may even be helpful in reducing neural noise, stabilizing the neural reward criteria (Cunningham et al., 2011; Koralek et al., 2012). The auditory feedback frequency could be erratic and quick to increase (potentially due to the hardware implementation), however, and the mice may not have

utilized the continuous feedback effectively. Therefore we wonder if the timescale of feedback was effective for learning in either of our studies.

The "Two-Learner" System

Brain-machine interfaces (**BMI**s) have shown remarkable progress in the last decade (Taylor et al., 2002; Carmena et al., 2003; Guenther et al., 2009; Gilja et al., 2012), and continued progress will one day help paralyzed patients regain lost motor function (Collinger et al., 2013). However, performance remains unacceptably low, considering the risks of invasive brain surgery. Increasing BMI performance partially relies on subjects learning how to use decoders, and partially on decoders adapting to changing neural activity of subjects. It is important to understand the roles that subject learning (Carmena et al., 2003; Wander et al., 2013) and decoder adaptation (Orsborn et al., 2012; Bishop et al., 2014) play in BMI skill acquisition, as well as how they interact (Taylor et al., 2002; Sanchez et al., 2009; Orsborn et al., 2014). Studies in which decoders are trained from direct motor observations show that learning to use the decoder is still a requirement of high performance (Taylor et al., 2002; Ganguly et al., 2011). This also suggests that decoder adaptation may play a part in the learning process (Vidaurre et al., 2011).

A primary concern for brain-machine interfaces (BMI) is the development of adaptive closed loop decoding algorithms that account for co-learning between the brain and the machine (the "two-learner" system; Shenoy and Carmena, 2014). As the brain

learns a task, neural plasticity alters activity patterns; concurrently, decoding algorithms must both adapt to the new patterns and improve performance.

While we looked at both sides of the co-adaptation problem, the two studies described in Chapters 2 and 3 have few direct comparisons in a "two-learner" context. The purpose of the learning and adaptive studies was to understand the neural and behavioral correlates of learning a neuromodulation BMI, and determine the causes of poor performance in an online adaptive decoding BMI performed by monkeys in a delayed saccade task, respectively. The neuromodulation learning study used low frequency oscillations as an indicator of a global state and was only used as a crude "switch" during the control of a one-dimensional auditory cursor; the adaptation study used high frequency field potentials to determine the intended direction of a saccade during delay period activity. Integrating multiple LFP frequency bands into the same decoder may enable complex motor BMI control (Bundy et al., 2016). The decoding methods used are indicative of the complexity of the respective tasks: the first study only used one logistic function to determine the frequency of an auditory tone; the second study used a recurrent neural network (composed of hundreds of logistical units) to determine intended saccadic direction. How a subject would learn to control the output of one logistic function, versus the recurrent output of hundreds of logistic functions is unclear. It is true that the first task was explicitly created for subject learning, for which the relatively simple decoder may be beneficial, and the second task was designed to minimize subject learning and use decoder adaptation (i.e. machine learning) in lieu of subject learning—which is not to say that subject learning could not co-exist.

Various bands in the beta and gamma range show coherence between peripheral muscles and sensory cortex, suggesting their importance for sensory processing (Baker et al., 2006; Witham et al., 2007). This fact is supported by the fact that neuromodulation of these and other bands affects the ability to process incoming information and to perform actions (Boulay et al., 2011; Schafer and Moore, 2011; Joundi et al., 2012; Khanna and Carmena, 2015; McFarland et al., 2015). Understanding the behaviors and underlying neural processes of neuromodulation is therefore an important step to take in the creation of competent BMIs. Problematically, if only specific beta rhythms like the tonic-hold-beta can be intentionally modulated, then it is possible that sensory-deafferented patient populations may never be able to perform this neuromodulation task as described (Kilner et al., 2004; Witham et al., 2007).

Fixed decoders may not be able to accommodate trial-trial variability, which an adaptive decoder might. The results from Orsborn et al. (2011) indicate that while trial success remained statistically similar before and after adaptation ceased, movement trajectories significantly declined once the decoder was fixed (Orsborn et al., 2011). This may be because fixed decoders are initialized from a limited repertoire of intended actions and thereafter shape subject neural activity to this initialization, causing the decoder to overfit the data and likely limiting the potential dimensions of control (Ganguly and Carmena, 2009). Wu and Hatsopoulos (2008) adapted a Kalman filter after each trial, which increased long term performance more than non-adaptive decoding (Wu and Hatsopoulos, 2008). This was likely due to unstable, drifting neural activity patterns.

These studies show the necessity of adaptive decoders that make reasonable assumptions about desired outcomes and are more amenable to exploratory learning (Orsborn et al., 2012; Fan et al., 2014).

It is vital to understand neurofeedback effects on neural plasticity, although decoder adaptation can mitigate alterations to the normal neural-behavior relationship, especially if appropriate situational context and goals are considered (Fan et al., 2014). Specifically in the context of the first study, if beta rebound is functionally associated with the maintenance of forward motor programs (Tan et al., 2014, 2016; Torrecillos et al., 2015; Cao and Hu, 2016) or are generated by tactile feedback (Parkkonen et al., 2015), then the specific selection of beta rhythm neuromodulation could be at odds with decoding algorithms' interpretation of increased beta power. Decoders that are not biologically informed may have trouble distinguishing between neural contexts, such as between preparatory activity and activity underlying desired action (Canolty et al., 2012). Adaptive decoders with more advanced knowledge of network interactions and brain function may better interpret intended outcomes (Kao et al., 2015; Panzeri et al., 2016). Additionally, subjects modulating any rhythm in a context not normally associated with the rhythm may reinforce undesirable cortical interactions.

A combination of learning and adaptation studies is thought to be a key step in improving BMI performance (Gilja et al., 2011). Before motor BMIs can be translated from the laboratory, it is critical for researchers to understand more of the dynamic between learning and adaptation (Ryu and Shenoy, 2009). Only when we can avoid

detrimental interactions between learning and adaptation, will their full potential as medical devices be realized.

APPENDIX

Learning Study Supplemental Information

Histology

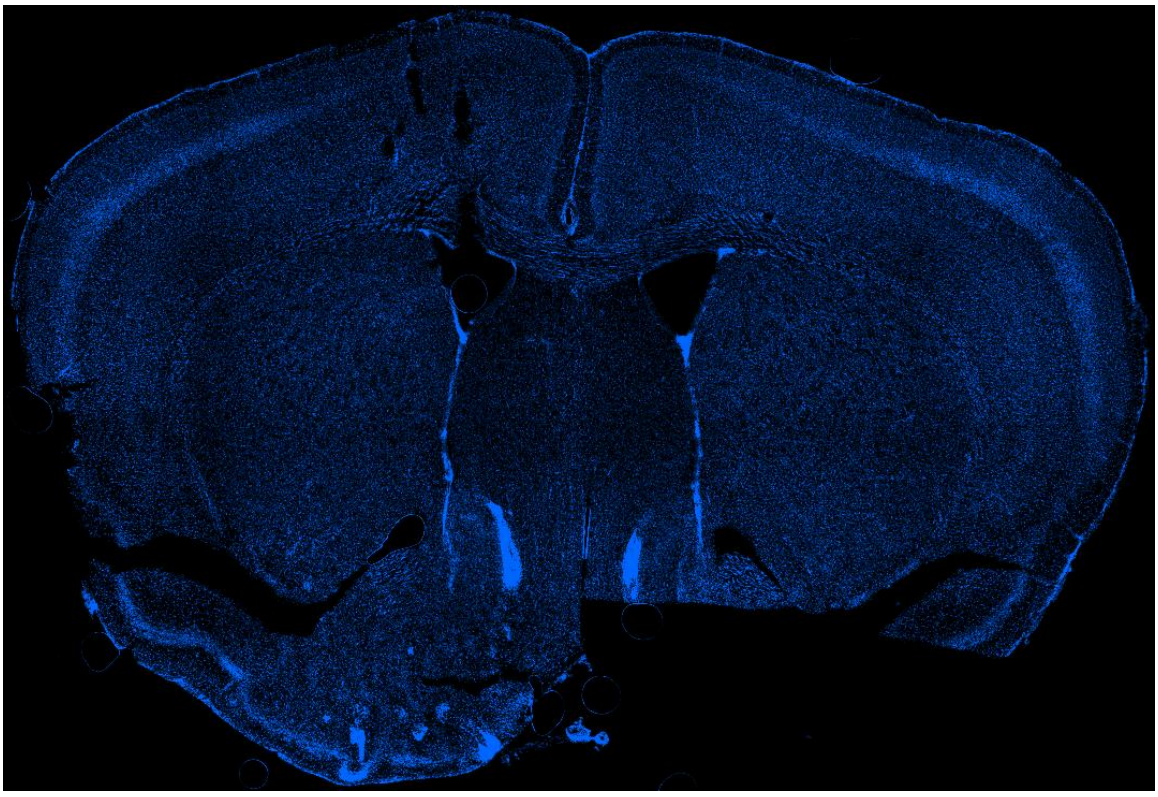

Figure 38: Mouse A histology

This slice shows the site of the reference electrode. Note the electrode tracts visible in the medial left hemisphere, in the deeper layers of vM1. The reference electrode was likely in the white matter or lateral ventricle (Paxinos and Franklin, 2001). This slice is located approximately +0.74 mm from Bregma.

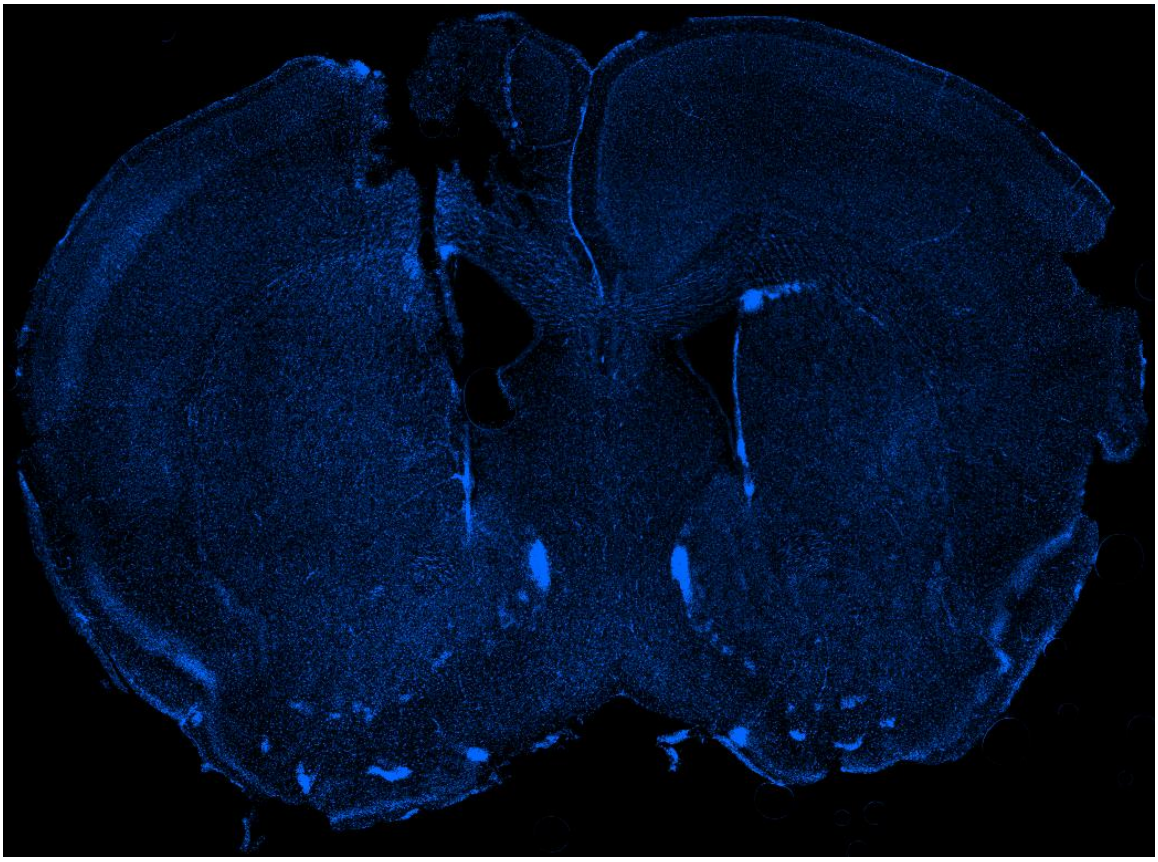

Figure 39: Mouse B histology

This slice shows the site of the reference electrode. Note the electrode tracts visible in the medial left hemisphere, in the deeper layers of vM1. The reference electrode tract is clearly visible to the left of the lateral ventricle, potentially in the lateral ventricle or striatum (Paxinos and Franklin, 2001). This slice is located approximately +0.86 mm from Bregma. There is more extensive damage to the cortex in this mouse as compared with mouse A, but because the mouse was sacrificed more than 6 months after surgery, it is unclear the extent of the damage during the relevant brain control sessions. Given the relatively good signals and

MUA recorded, we suspect the damage may have been minimal during that time.

"Bad" Electrodes

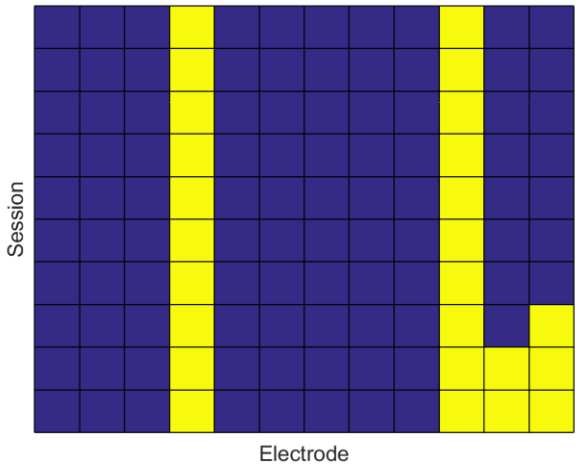

Figure 40: "Bad" electrodes

"Bad" electrodes (in yellow), as determined by visual and auditory inspection at the start of daily sessions for mouse C. These electrodes were not included in the online BPR calculation, and are ignored in all data analyses. Note that session 1 is the top row and session 10 is the bottom row.

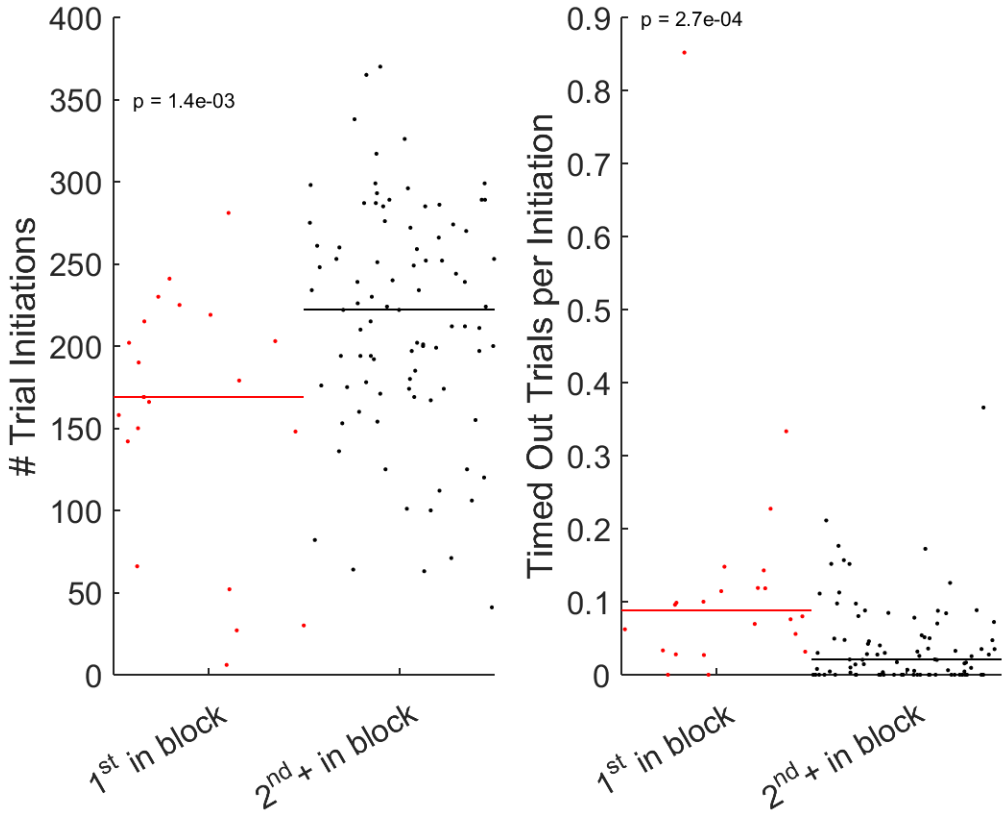

Figure 41: Low motivation for reward within 24 hours of deprivation

Left: Mice initiate significantly fewer trials during the 1st session of a contiguous block of sessions than during the 2nd and later sessions within the block (median 169 vs. 222, $p = 0.0014$, Wilcoxon rank sum test, Holm-Bonferroni corrected). Right: Significantly more timed out trials also occur per trial initiation during the 1st session of a contiguous block (median 0.09 vs. 0.02, $p < 0.001$, Wilcoxon rank sum test, Holm-Bonferroni corrected).

BPR-Auditory Tone Transformation

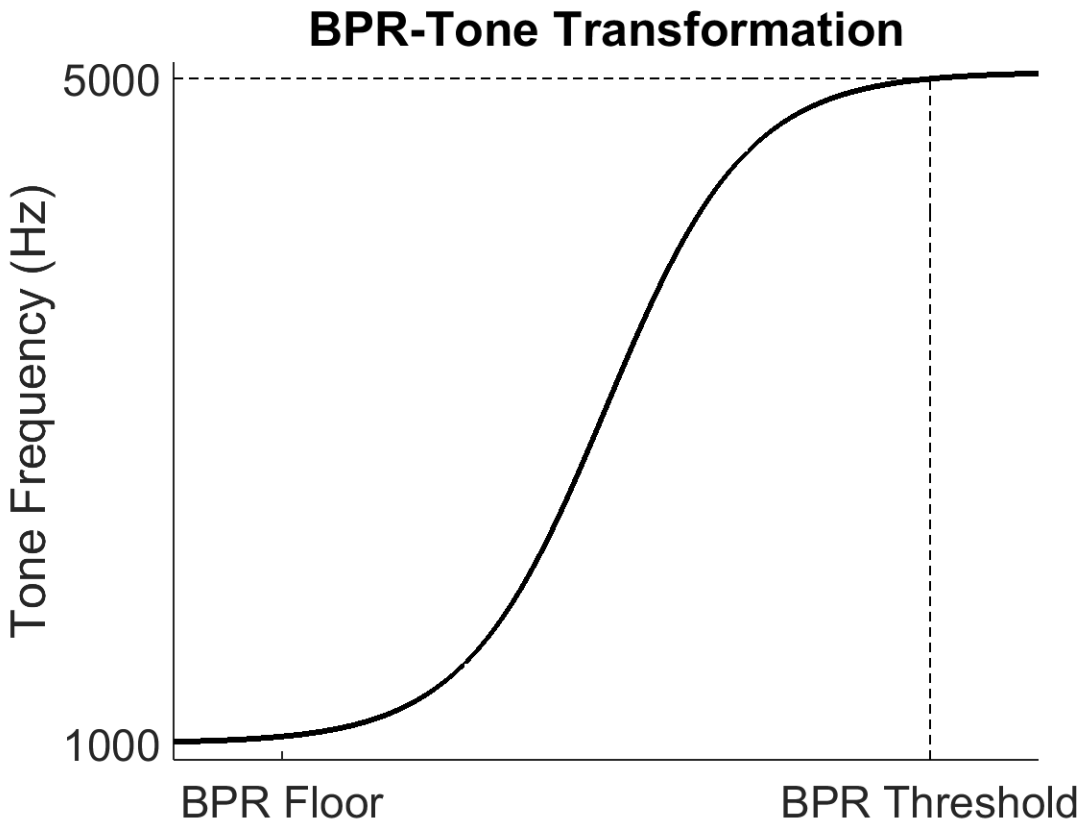

Figure 42: BPR-auditory tone transformation

We found a logistic transformation to be the most appropriate form to use, given that the BPR range could vary slightly between sessions, and the range-limit of the logistic curve reduced the need for parameter changes.

MUA Quality Control

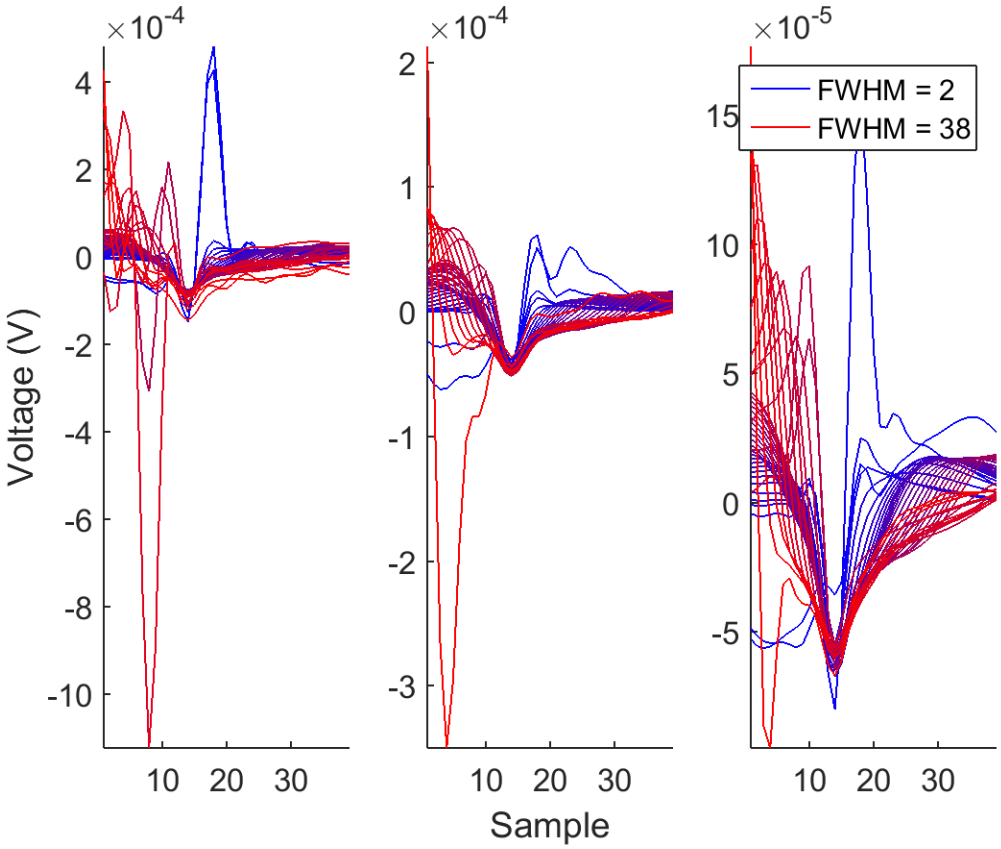

Figure 43: Mean MUA shape grouped by FWHM

The full-width half-max value color coding shifts from blue (FWHM=2) to red (FWHM=38). The reddish waveforms tend to display non-neural waveform characteristics, while the blueish waveforms tend to look more neural. The three columns show mean MUA activity for mice A, B, and C, respectively.

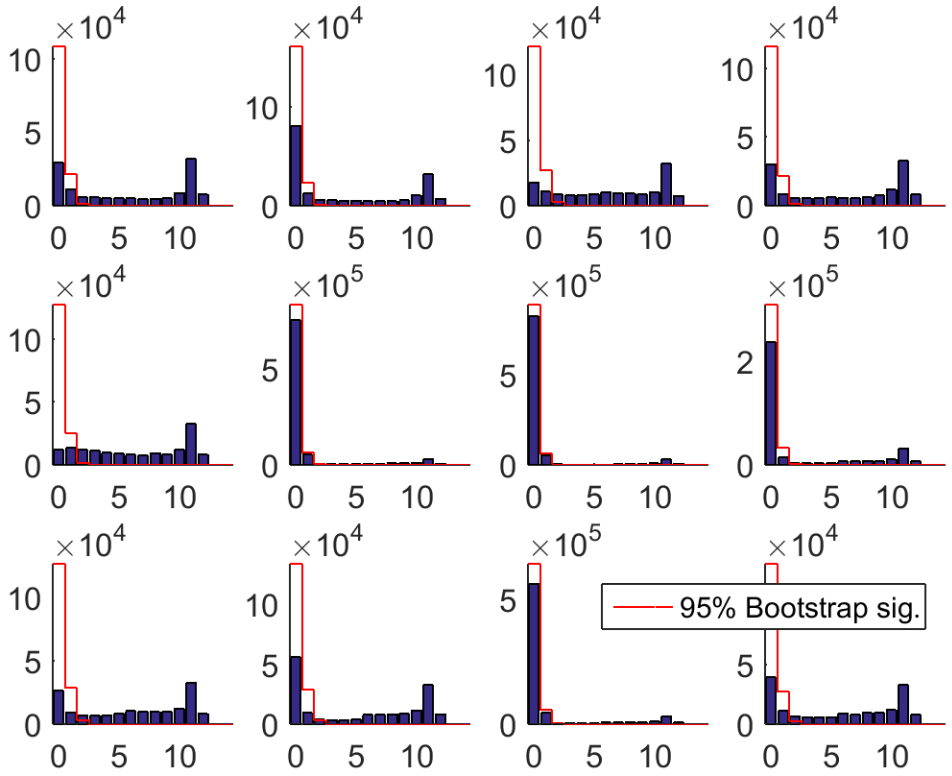

Figure 44: MUA-coincidences based on chance

Shown is one mouse. We attempted to determine the frequency with which MUA activity should co-occur on 12 different electrodes, binned in 1 ms samples. Based on the 95% bootstrapped confidence intervals, it appears that no more than 2-3 electrodes should have had coincident MUA activity in the same 1 ms interval. MUA activity co-occurring on more than ~3 electrodes is increasingly likely to be due to noise or motion artifacts. To eliminate erroneous MUA activity, without removing coincident MUA activity due to common neural activity, we excluded all MUA activity that co-occurred on 6 or more electrodes.

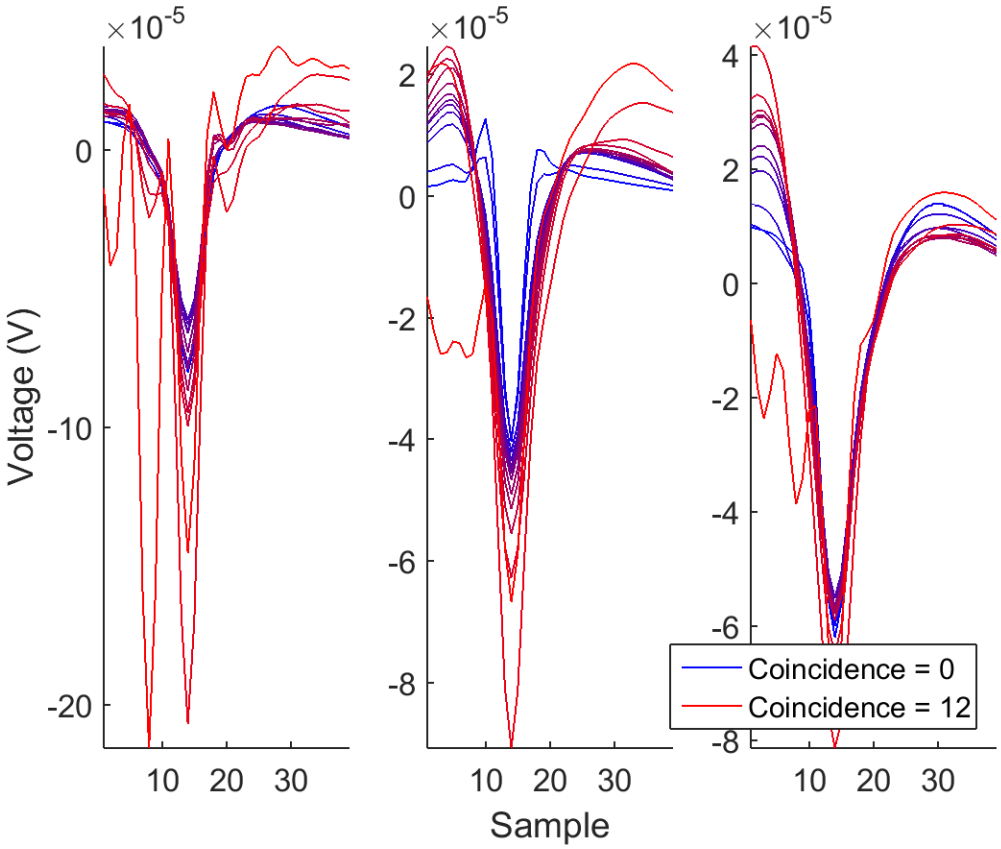

Figure 45: Mean MUA shape grouped by coincidences

Shown are the $\boldsymbol{n}$ coincidences for all three mice. The color coding shifts from blue (no other coincidences) to red (12 coincidences). Highly coincident MUA activity tends to display non-neural waveform characteristics. Note that it was possible to have more than 11 coincidences if multiple spikes occurred within the same 1 ms bin on the same electrode (note the double peaks for the red traces representing 12 coincidences). The three columns show mean MUA activity for mice A, B, and C, respectively.

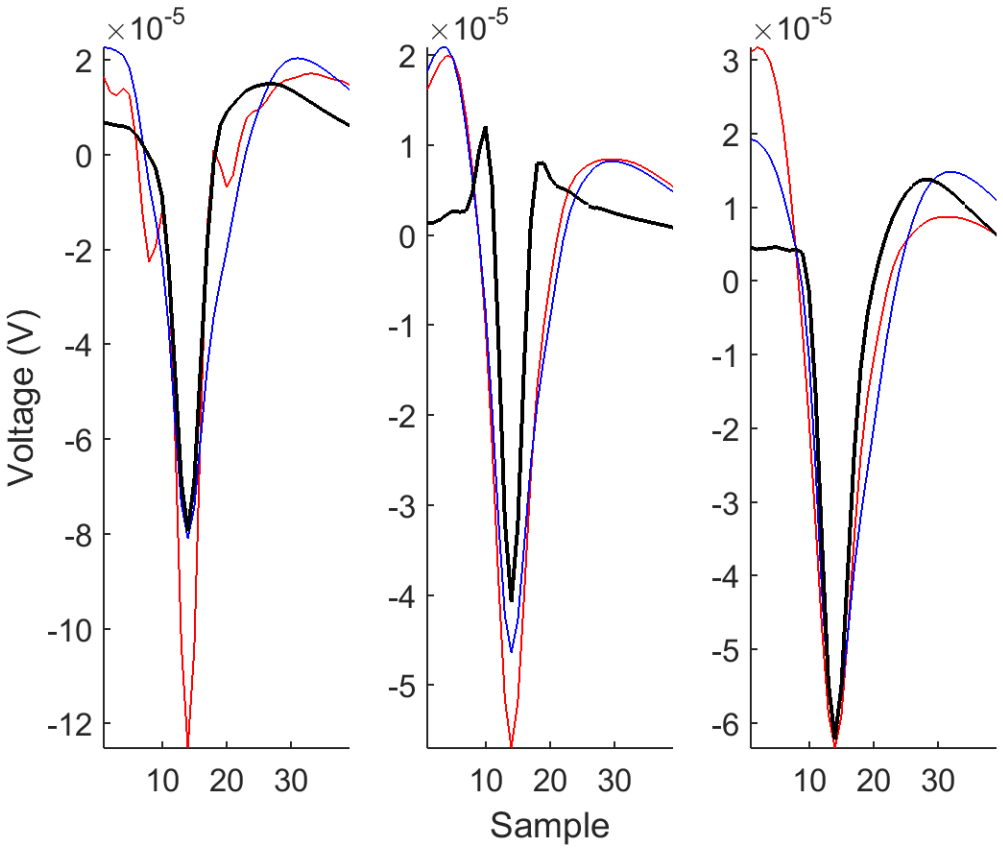

Figure 46: Mean quality-controlled MUA shapes

In black are the final mean MUA shapes over all electrodes. These MUA appear neural and were used in all MUA-based analyses. The three columns show mean MUA activity for mice A, B, and C, respectively.

Video Analysis

Slow speed video (480 x 720 pixels/frame, 29.97 frames/s) was used to provide measures of behavioral activity. Frame alignment was necessary for multi-session and within-session video analysis, due to non-fixed positions of the camera and the presence of jitter. Video frame alignment was performed in three steps:

(1) Due to the presence of jitter, each video frame potentially needed an X- or Y-shift correction to an intra-session common average frame. The common average frame was calculated via the following algorithm:

STEP 1: For all video frames *i*,

$$XYShiftCorrections_i = \text{argmin}_{x,y}\, xcorr2d(frame_i, \langle frame \rangle)$$

STEP 2: $\langle frame \rangle = \frac{1}{n}\sum_n (frame_i + XYShiftCorrections_i)$

STEP 3: REPEAT STEPS 1 AND 2 IF $\Delta XYShiftCorrections_i > \varepsilon$

(2) Due to the fact that the position and perspective of the video camera was not identical between sessions, an affine transformation was calculated for each jitter-corrected intra-session mean video frame (see (1)). The mean jitter-corrected frames from each session were annotated to mark the four corners of the behavioral arena, and an affine transformation was calculated from these four corners and the inter-session mean of the four corners for all sessions.

(3) Each individual video frame from a given session was corrected by applying the de-jittering X-Y-shift (1) and affine transformation (2).

The results of the first two steps of the alignment process can be seen in Figure 47, in which we can see the inter-session mean frame of the mean frames across all sessions from mouse C.

Figure 47: Video frame alignment across sessions

Left: the inter-session mean video frame before affine transformation correction, consisting of 10 intra-session-averaged video frames from mouse C. Note the appearance of multiple shifts of the LED (lower left) and the experimental beam break ports (right), indicating that the video camera was not in the identical position from session to session. Right: the inter-session mean video frame after affine transformation correction. Note that there is much less "ghosting" or shifting than the pre-corrected image, and that the experimental beam break ports (at right) are clearly visible. Also note that the LED (lower left) shows ghosting and shifts, due to the fact that the LED was frequently adjusted between sessions. Also note the blurry figure in front of the two behavioral ports, on the right, which is where the mouse was most commonly situated.

Electrode Correlations

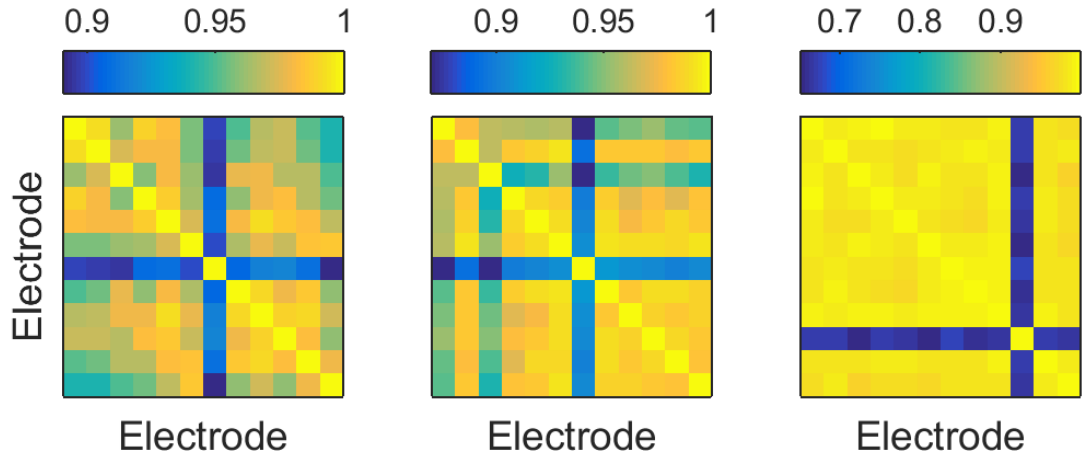

Figure 48: Electrode correlations around BPR threshold

Electrode 7 (for both mice A and B) were not marked as "bad" during the recording sessions and it is unclear why their correlation is relatively lower (~0.9) than the other electrodes. Electrode 10 (for mouse C) was marked as "bad" during all recording sessions, so its relatively low correlation (~0.7) is not unexpected.

FA/TP Latency within Sessions

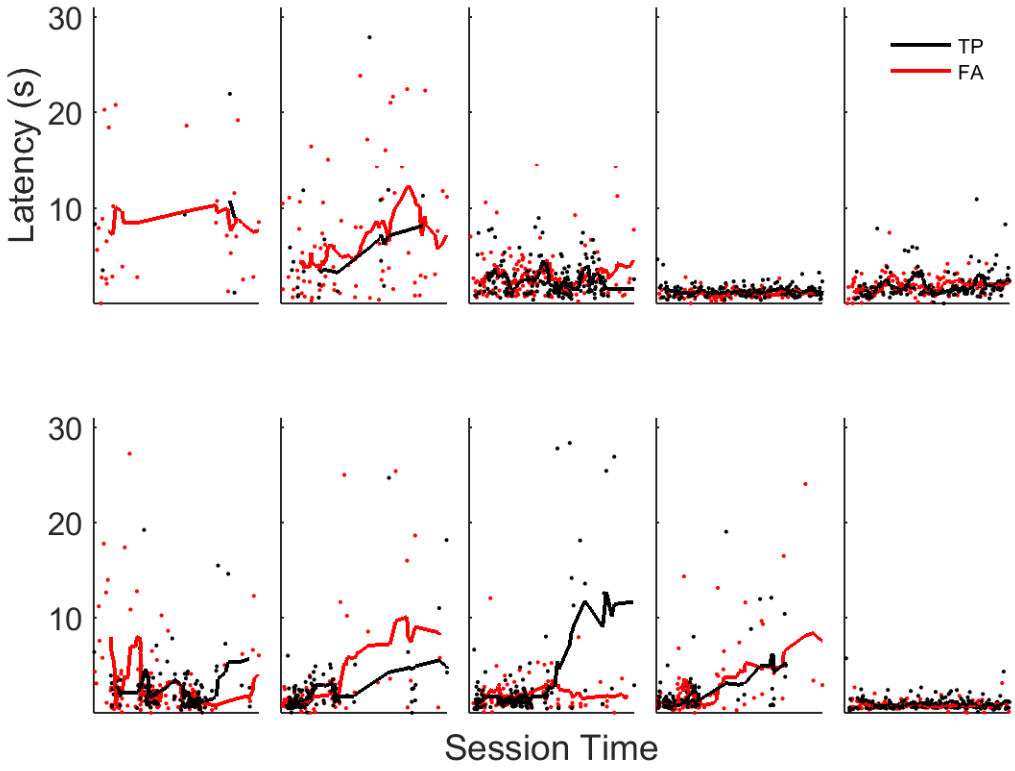

Figure 49: FA/TP latency for mouse A

Red: 10-pt moving average of FAs. Black: 10-pt moving average of TPs.

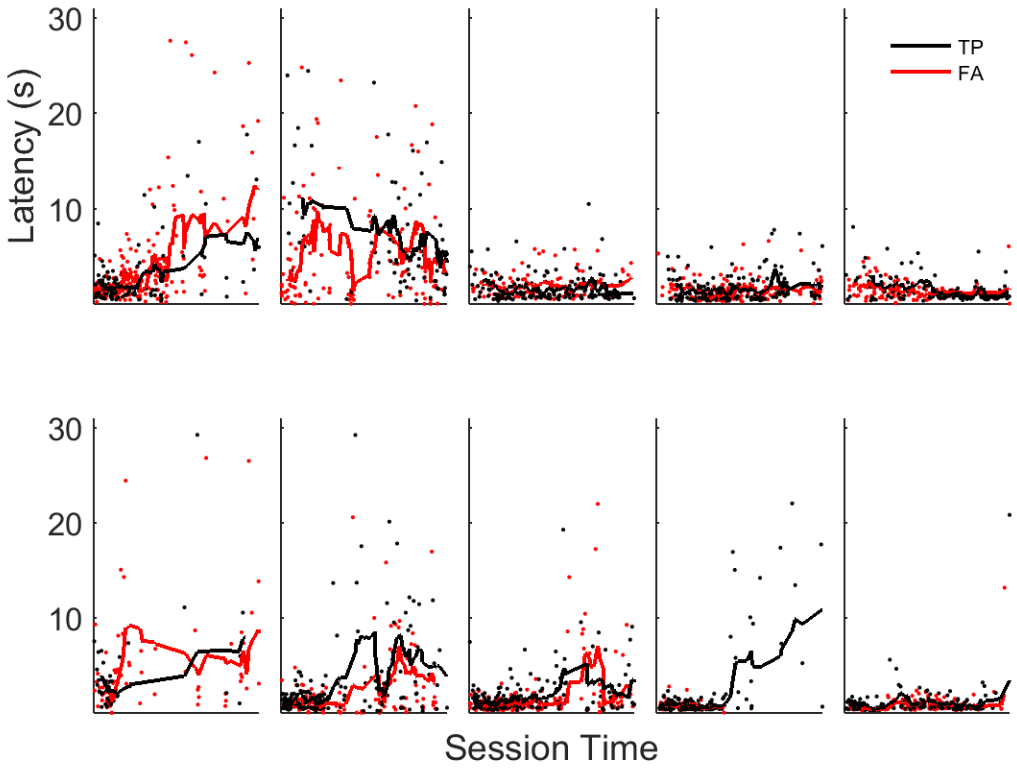

Figure 50: FA/TP latency for mouse B

Red: 10-pt moving average of FAs. Black: 10-pt moving average of TPs.

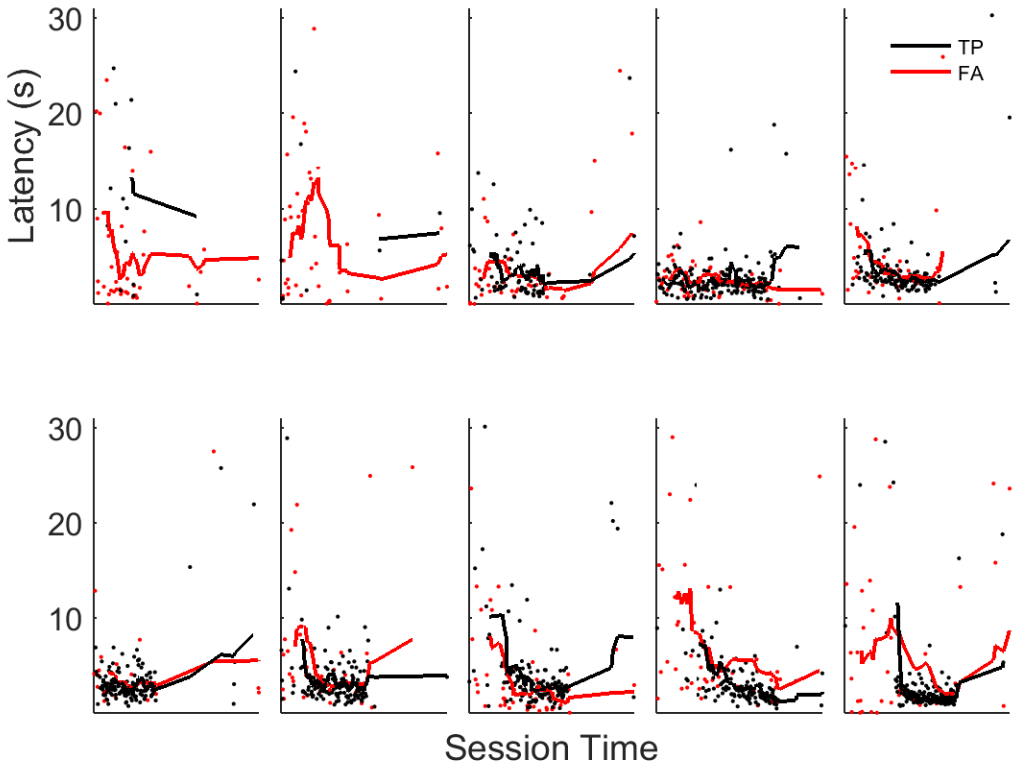

Figure 51: FA/TP latency for mouse C

Red: 10-pt moving average of FAs. Black: 10-pt moving average of TPs.

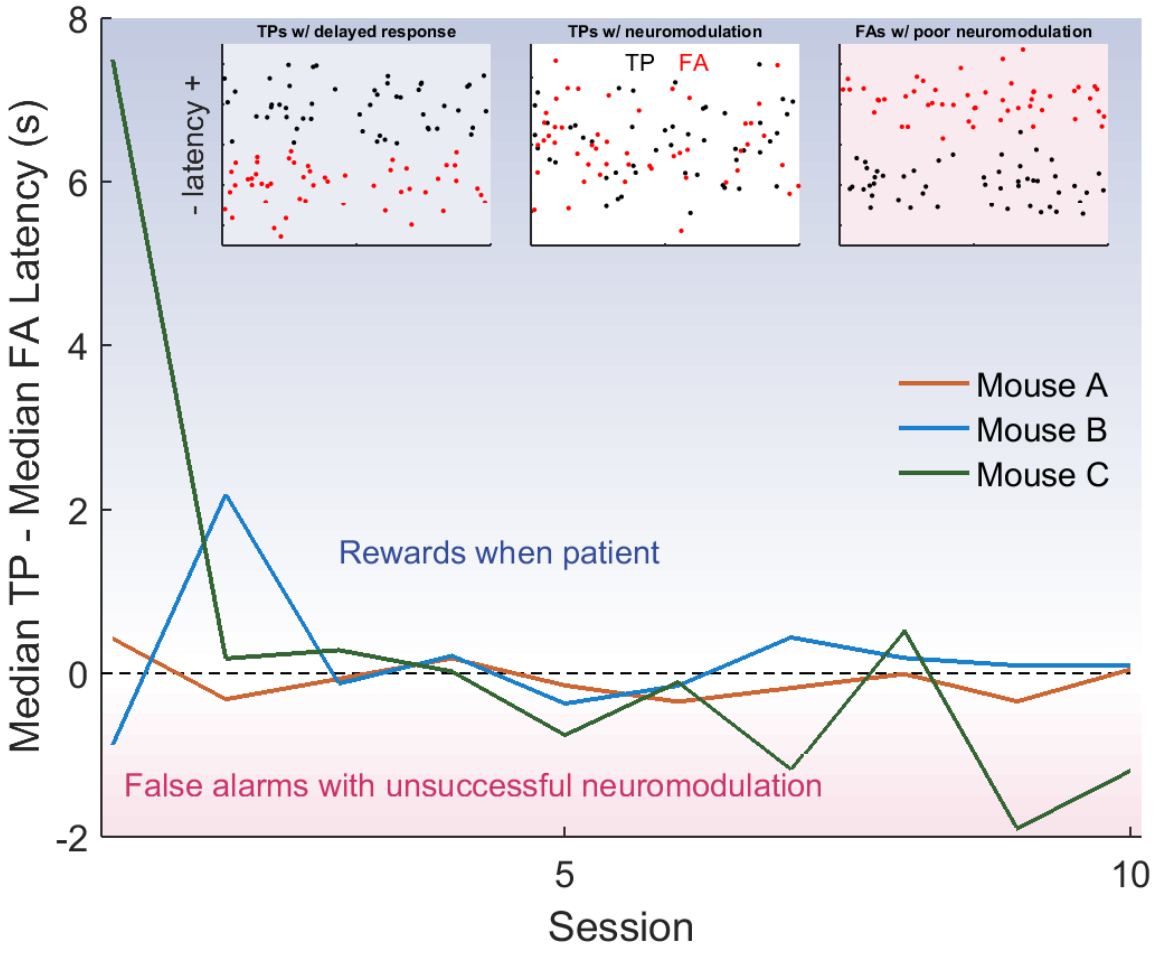

Figure 52: Behavioral Index increase not due to behavior

Lack of difference between median true positive and false alarm latencies indicates random reward seeking and suggests behavioral index improvement is due to neuromodulation performance. Inset, top left: longer true positive latencies than false alarm latencies may indicate that reward status is due to delayed reward seeking response (i.e. if mice wait long enough, there is a growing chance that reward will be available). Inset, top middle: similar true positive and false alarm latencies may indicate that reward status is due to neuromodulation performance (i.e. if

mice seek reward randomly, then earlier reward availability increases the chances of success). Inset, top right: longer false alarm latencies than true positive latencies may indicate that false alarms are committed after the "usual" reward availability latency has passed (i.e. if mice wait long enough without reward, there is a higher chance they will seek reward). The first behavioral regime (TP latencies longer than FA latencies) is indicative of behavioral and/or neuromodulation inexperience, and the third behavioral regime (FA latencies longer than TP latencies) is indicative of behavioral and/or neuromodulation learning.

Trial Progressions Over Sessions

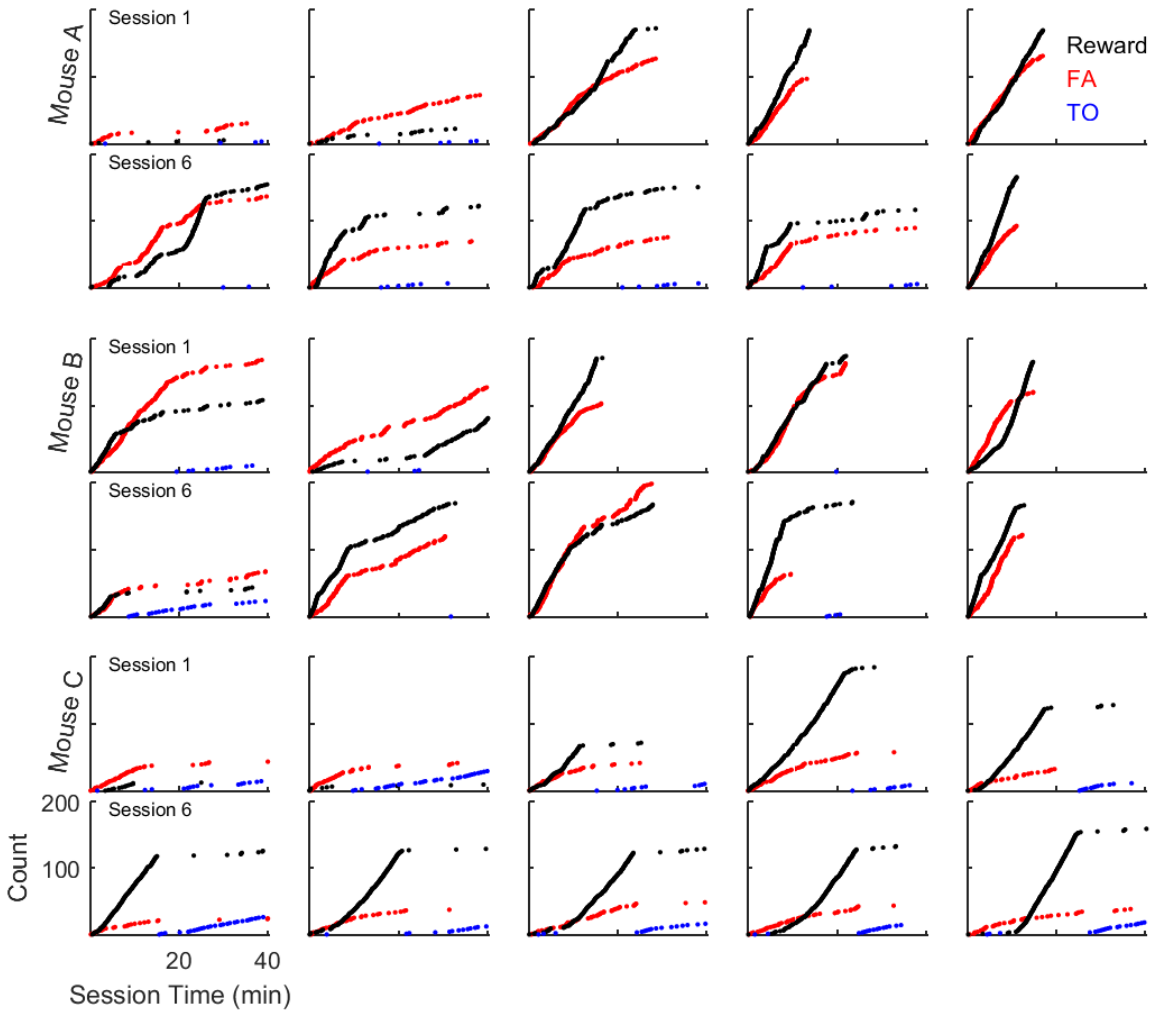

Figure 53: Distribution of TPs, FAs, and TOs within sessions

Distribution of true positives (rewards), false alarms, and trial time outs within sessions suggest the following: (1) mice committed more false alarms than true positives both during the first few sessions of the experiment and often during the beginning of later sessions (e.g. mouse C, session 10); and (2) trial time outs occur later in sessions, when mice are more satiated (e.g. mouse B, session 1). Events are plotted at the time at

which they occurred on the x-axis, and at the number of previous similar event types (i.e. rewards, false alarms, or trial time outs) on the y-axis. Sessions are plotted in order from left to right.

Mouse Response to Auditory Feedback

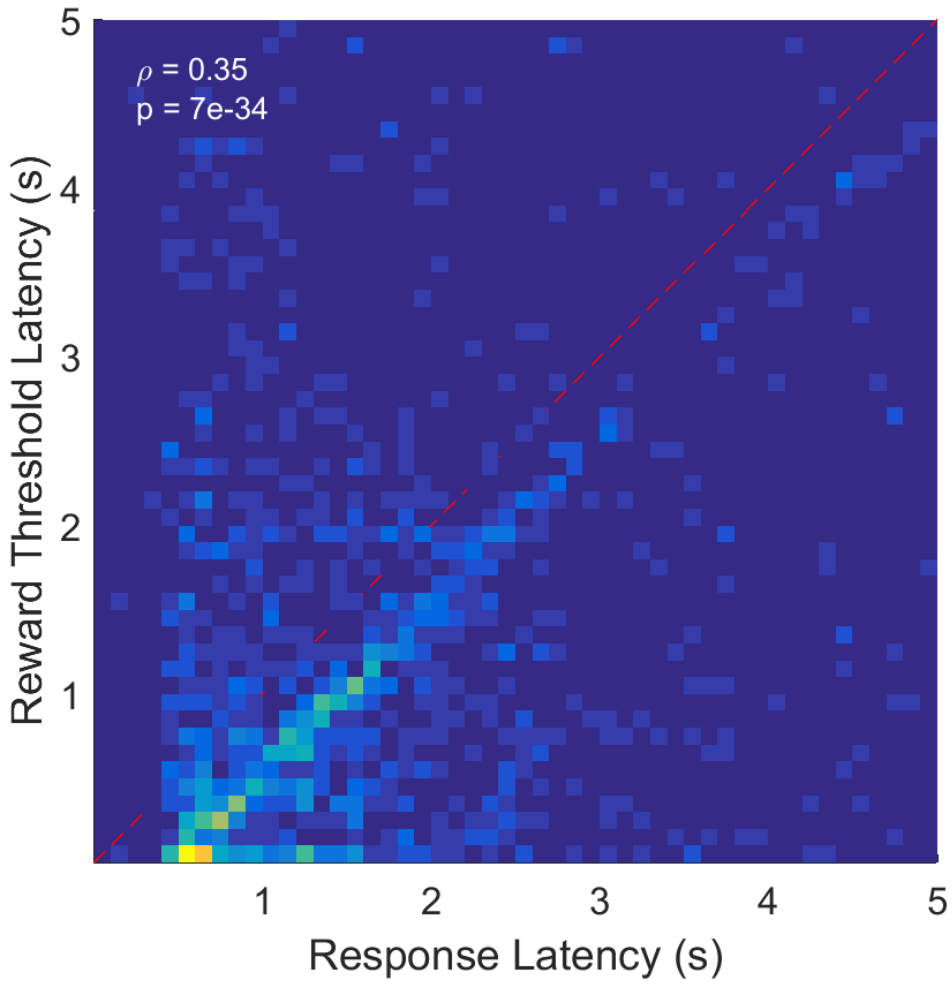

Figure 54: Mouse C seeks reward ~400 – 600 ms after reward tone

Mouse C performed several brain control sessions in which the feedback was pre-recorded—which therefore pre-determined the reward latency. From this data, we are able to reconstruct the actual distribution of reward

seeking responses to the true reward time. There was a significant correlation between reward threshold latency and response latency ($\boldsymbol{log_{10}\, p < -33}$; Spearman rank correlation), with a clear density of responses ~400 – 600 ms after the reward tone. Red dashed line: reward tone.

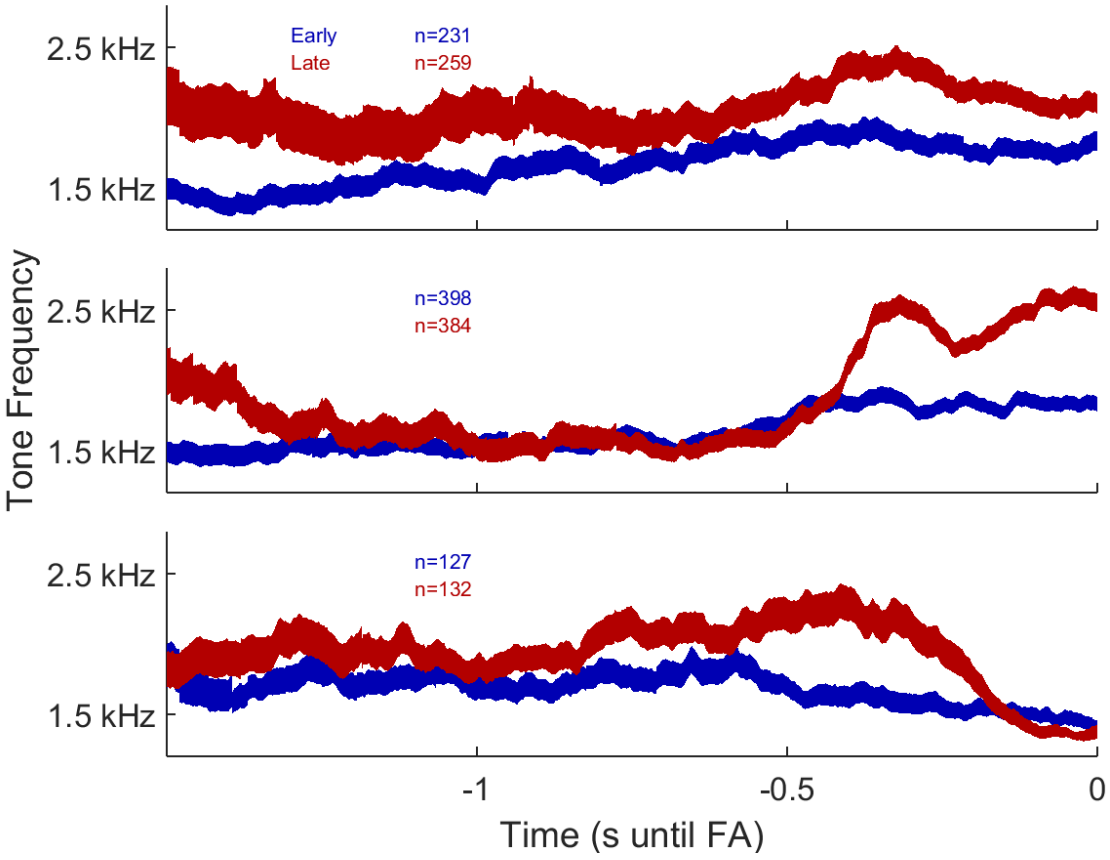

Figure 55: Auditory feedback frequency is possibly informative

Average auditory tone frequencies of all false alarm (FA) trials suggest utilization of frequency information in feedback. The 2nd and 3rd mice show some signs that they utilize the frequency content of the auditory

tone in their behavioral responses. FAs committed during late sessions (sessions 8-10; red) occur when the auditory frequency is higher than when FAs were committed during early sessions (sessions 1-3; blue). Middle: the late FAs from the 2nd mouse happen within ~0.5 s after a 1000 Hz increase in the auditory tone. Bottom: the auditory frequency around ~0.5 s before late FAs is higher than in early FAs. All traces are mean ± s.e.m.

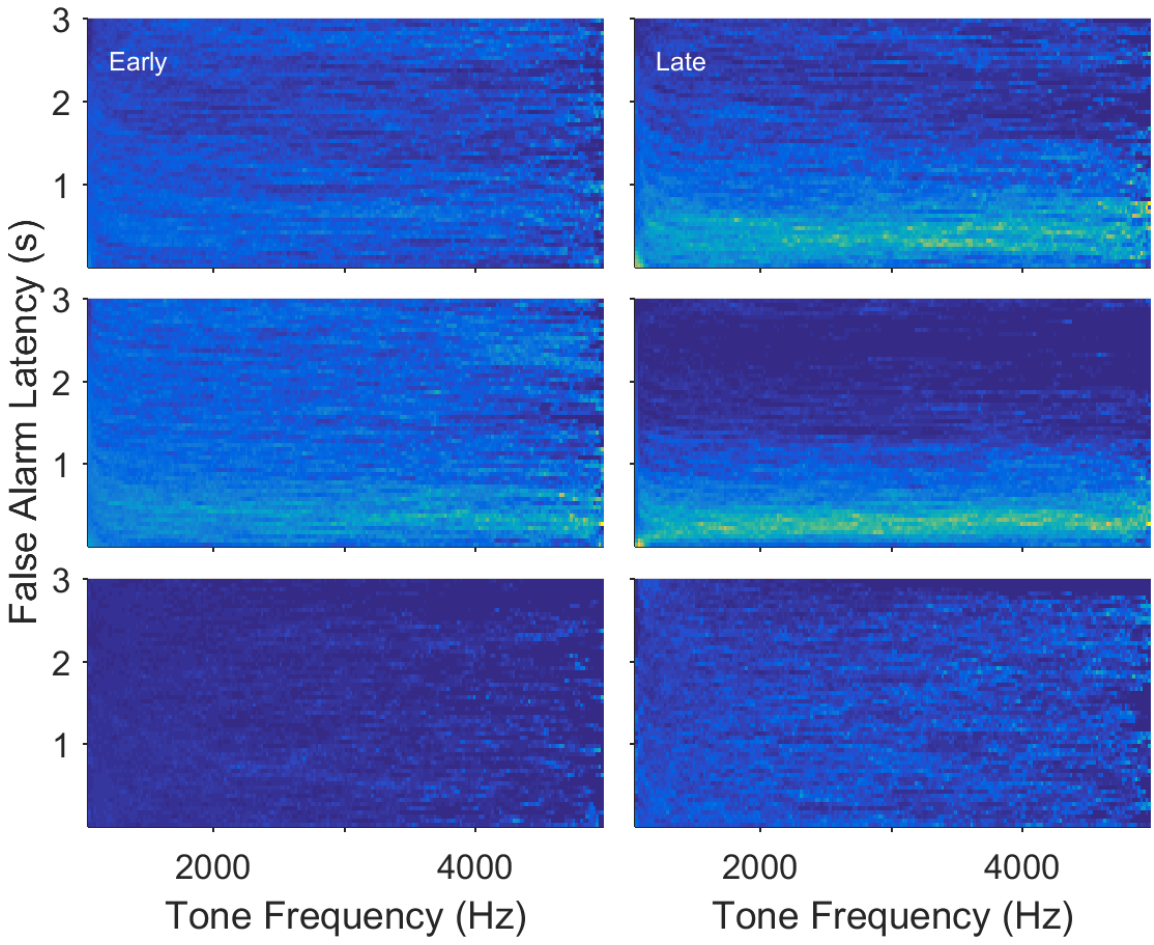

Figure 56: Little or no dependence on feedback frequency for reward seeking

Density of false alarm response latencies to tone frequency shows little or no dependence of reward seeking behavior on feedback frequency. Top row: mouse A shows a slight increase in reward seeking responses ~750 ms after 4.5+ kHz auditory frequency in late sessions (top right). Middle row: mouse B shows more high frequency-dependent reward seeking during early sessions (middle left) than in late sessions (middle right). Bottom row: mouse C shows no clear FA response behavior based on auditory frequency. All bins are normalized by the total number of events within each frequency bin.

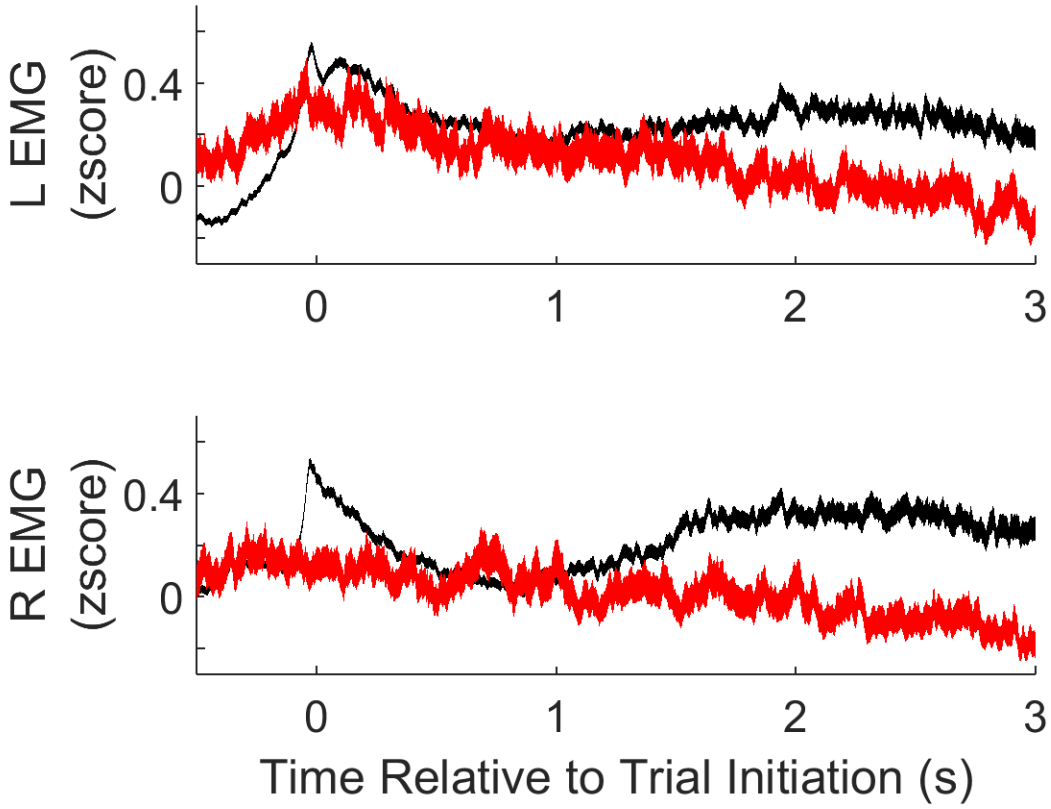

Figure 57: Increased EMG associated with trial success

Black: mean (± s.e.m.) EMG activity during true positive trials, up to, but not including, reward threshold. Red: mean (± s.e.m.) EMG activity during timed out trials.

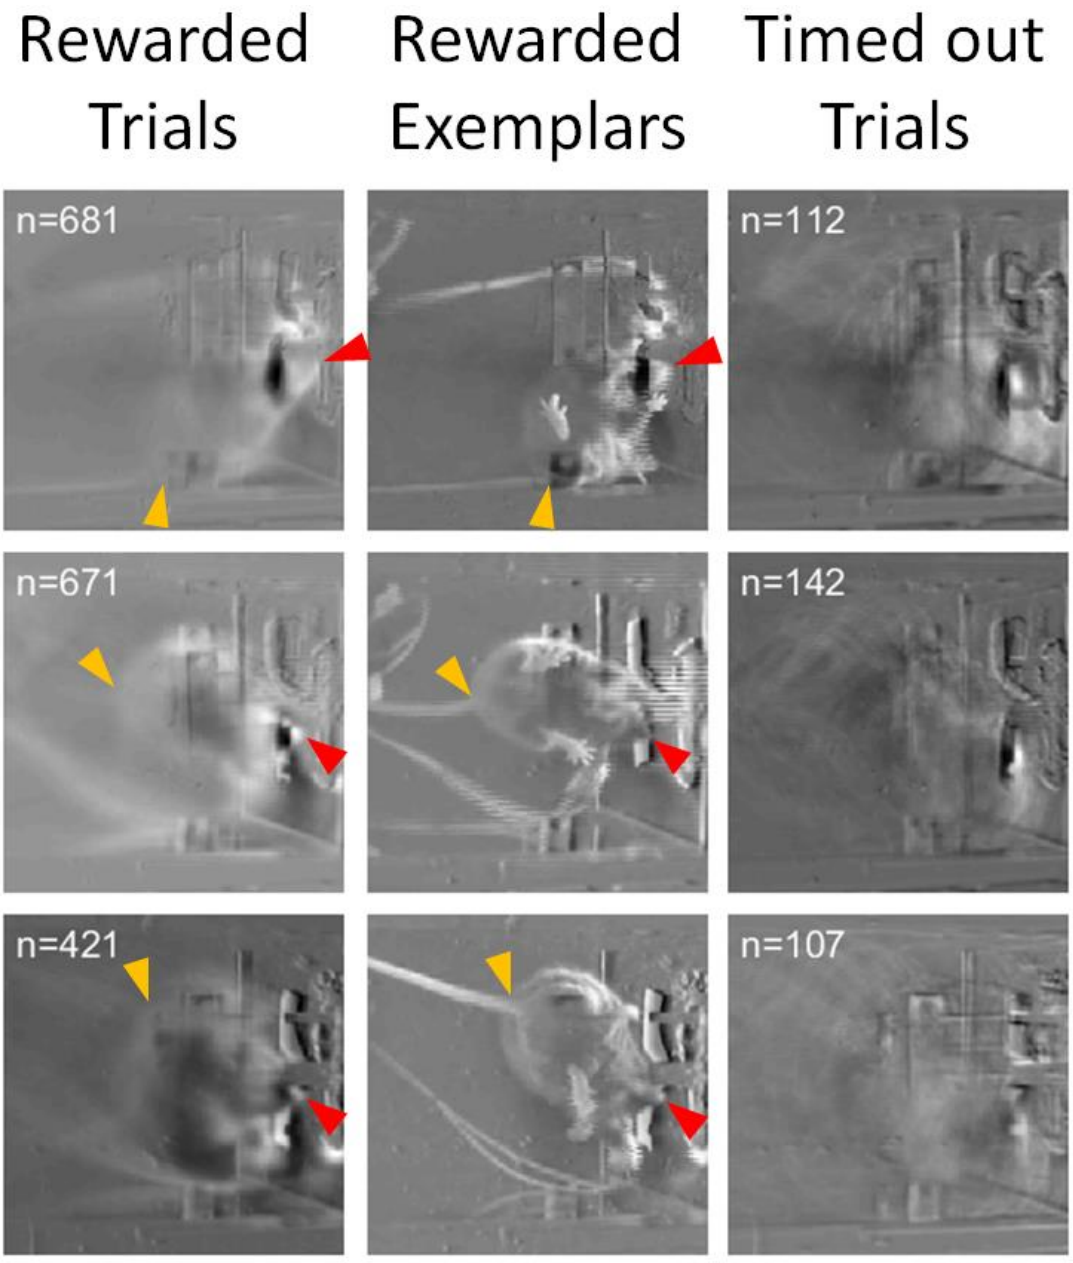

Figure 58: Mouse position at trial initiation predictive of outcome

Timed out trials occur when mice do not intend to participate in the neuromodulation task. Left column: mean frame at trial initiation for all true positive trials. Center column: single video frame with the highest correlation to the mean trial initiation frame (at left), for comparison.

Right column: mean frame at trial initiation for all timed out trials. Note that the general shape and features of the mice are visible in the frames in the left column, and that there is no discernible mouse shape in the frames in the right column, despite there being fewer frames in the average (e.g. n=112 vs. n=681 in the top row). Red triangles: position of the mouse snouts. Orange triangles: position of the mouse rumps, where the tail meets the abdomen.

LFP-EMG Correlations

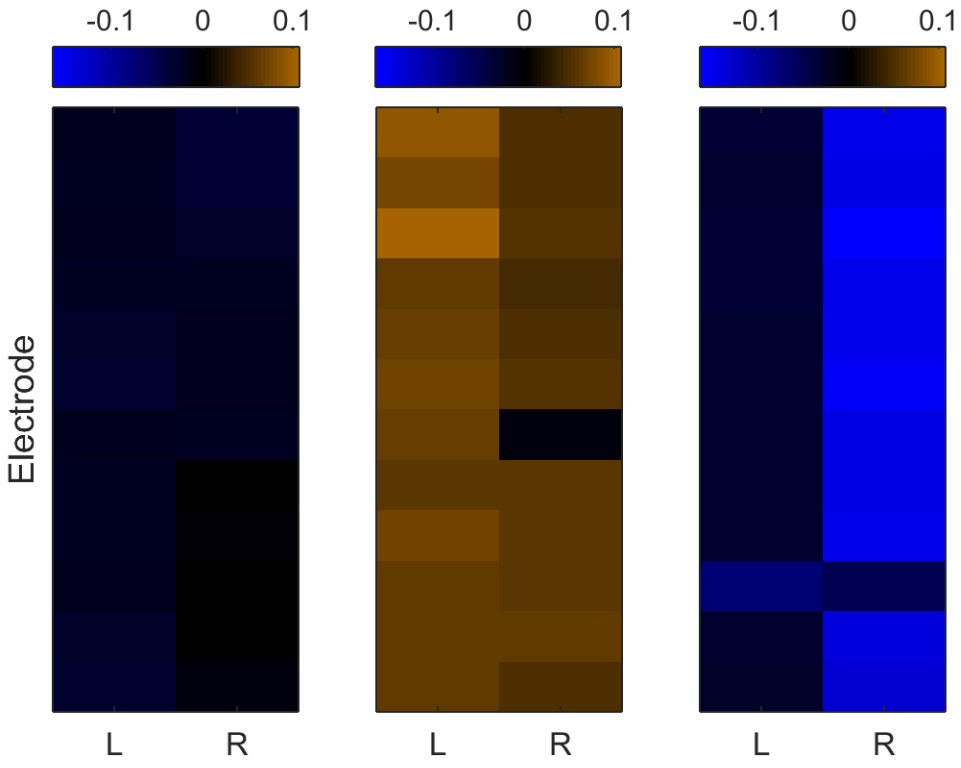

Figure 59: Low EMG-LFP correlation

Lack of EMG-LFP correlation suggests little or no motion artifact is present in the electrode activity.

LFP-EMG Coherence Phase

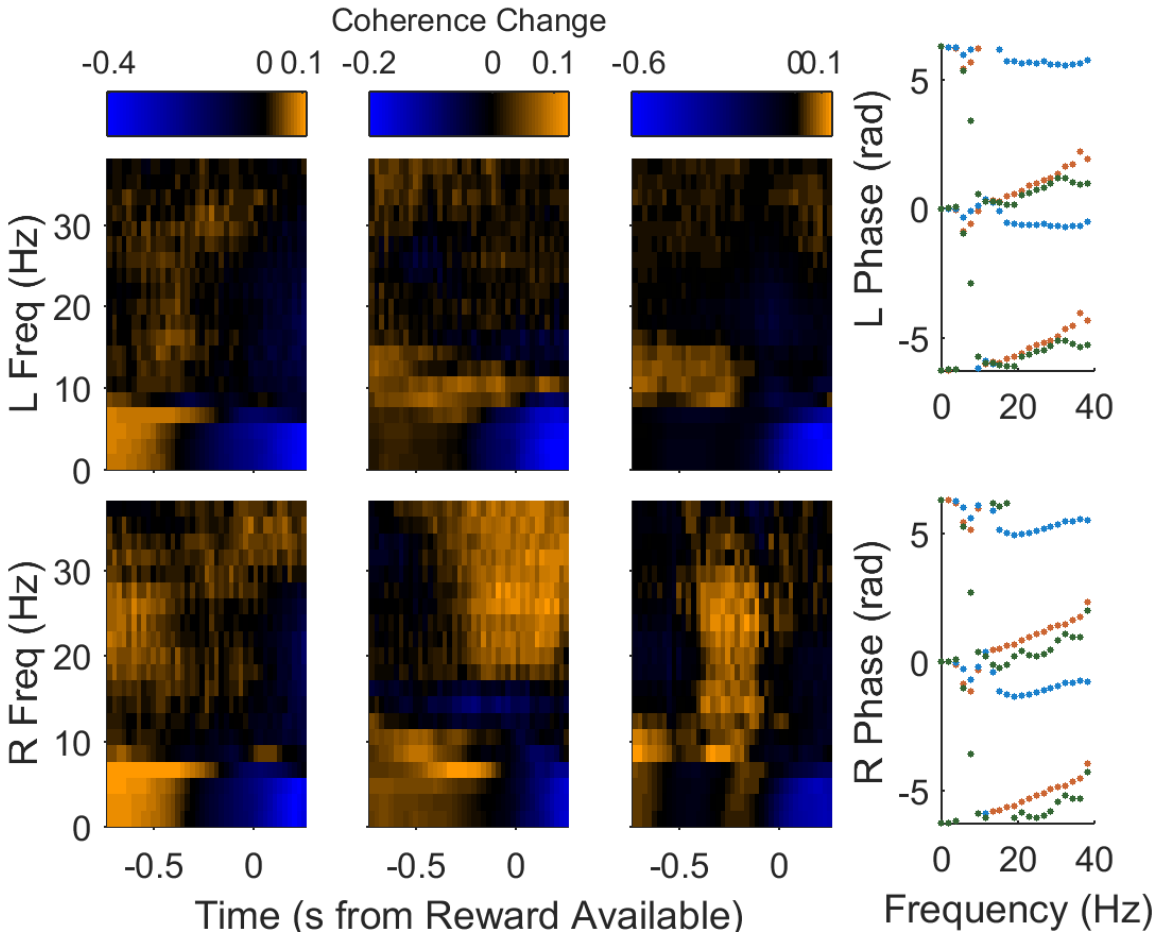

Figure 60: Increased contralateral EMG-LFP coherence around reward availability

More consistent increases in beta-range and theta-alpha-range right EMG-LFP (bottom row) coherence prior to reward availability than in left EMG-LFP coherence (top row). Both left and right EMG have an increase in ~6 – 12 Hz LFP-EMG coherence prior to reward availability, but right EMG-LFP show a consistent increase in coherence within beta-range

frequencies. The positive slopes of phase in the beta range (right column) for the right EMG suggest the vM1 LFPs lead EMG and coherence is driven by vM1, not afference from the whisker pads (Schoffelen et al., 2005).

MUA-EMG Phase Preference

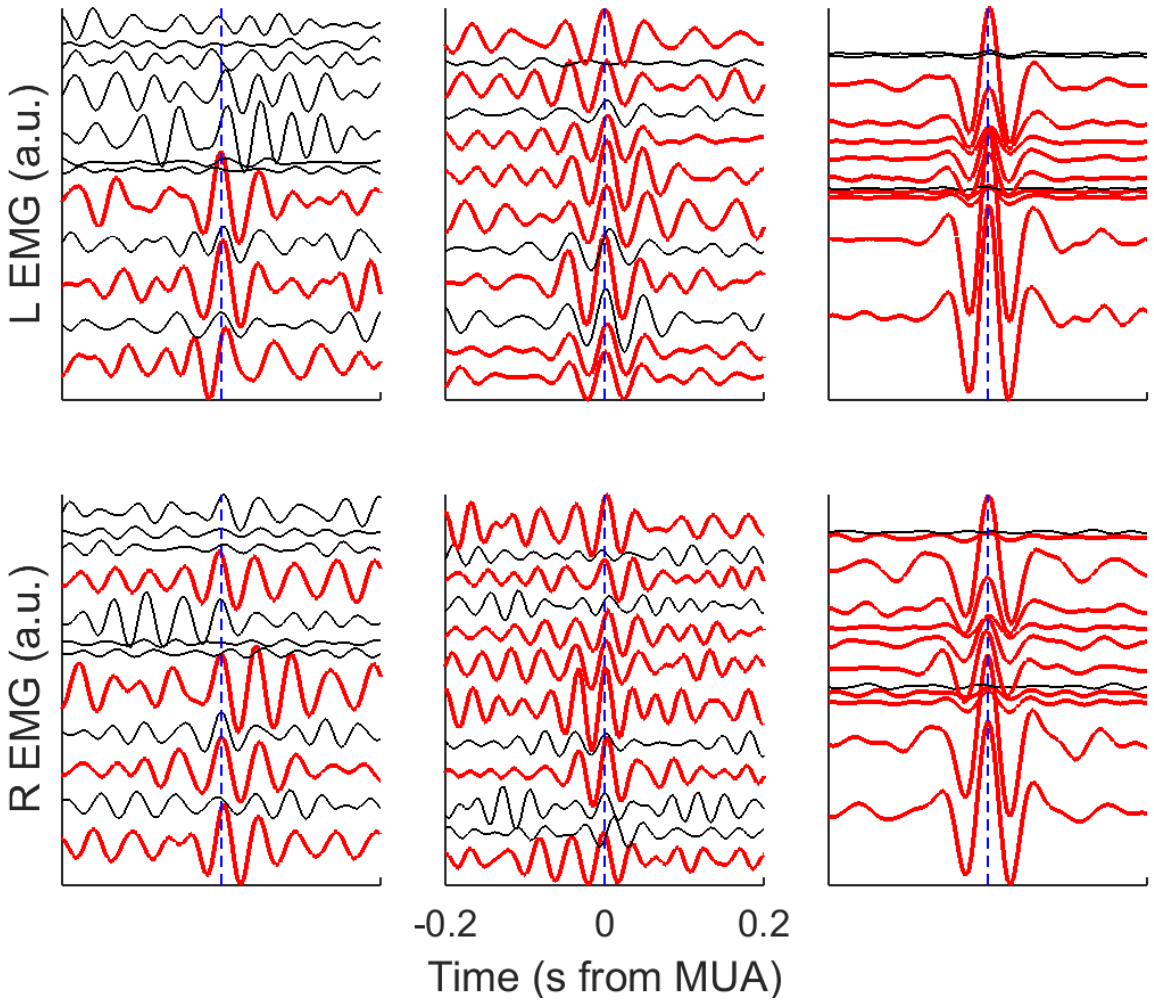

Figure 61: MUA-triggered EMG traces show protraction phase preference

MUA activity was found for each of the 12 electrodes, and then EMG traces were filtered (frequencies correspond to Fig. MUA Phase) and

averaged together. The peaks of EMG activity should correspond to the start of whisker protraction (Schroeder and Ritt, 2013). Red traces: mean EMG for phase-selective MUA. Black traces: mean EMG for non-phase-selective MUA.

EMG-EMG Coherence Changes

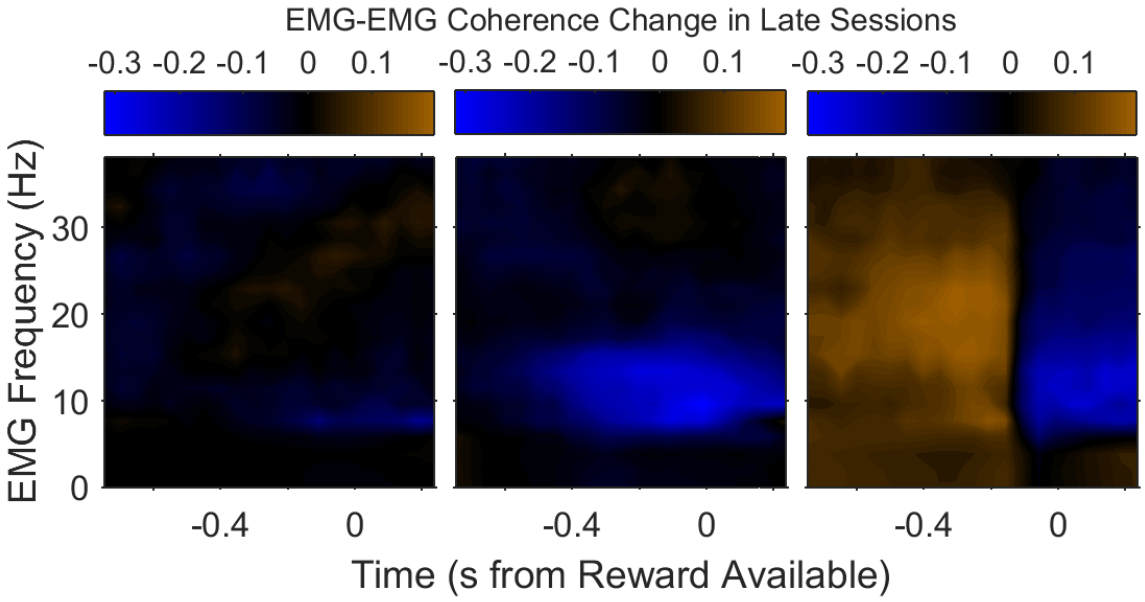

Figure 62: EMG-EMG coherence decreases around reward threshold

8 – 16 Hz EMG-EMG coherence decreases around reward threshold in late control sessions.

Burst Durations vs. Amplitudes

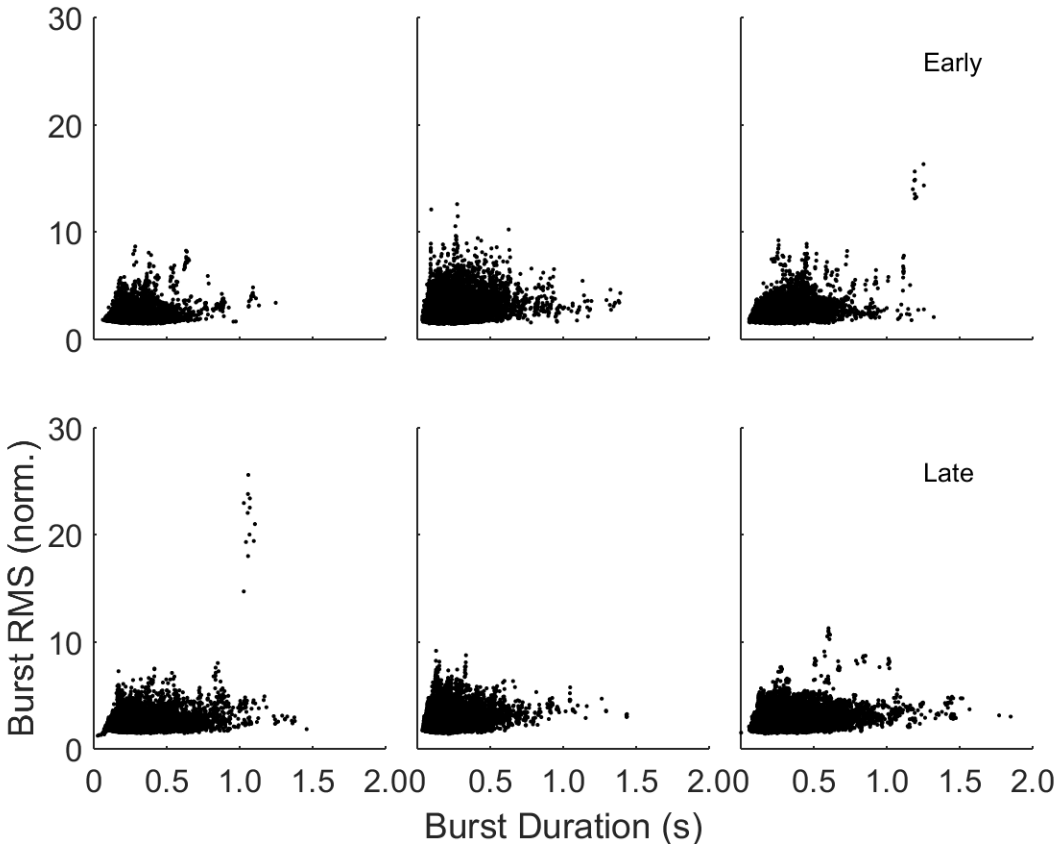

Figure 63: Burst duration vs. session-normalized power

Burst duration vs. session-normalized power in early (top row) and late sessions (bottom row). The duration of the bursts appears to have no relation to the power of the bursts, arguing against the idea that the increased duration is due to higher overall burst power.

Adaptation Study Supplemental Information

Batch Model Learning Parameters

Learning Rate	Momentum	L2 Weight Decay	Dropout Rate	Bayesian Targets	Input Resampling Rate
0.0025	0.95	0.0001	0.2	0.87	0.04

Table 1: Batch learning parameter values

Online Model Learning Parameters

Learning Rate	Momentum	L2 Weight Decay	Dropout Rate	Bayesian Targets	Input Resampling Rate
0.0001	0	1e-20	0.1	0.87	0.01

Table 2: Online learning parameter values

Generalization Techniques

Dropout approximates the training of 2^n different network architectures (where n is the total number of hidden units) by zeroing a random subset of the hidden units' outputs during each training iteration (Srivastava et al., 2014). For each trial in a training batch, a hidden unit's activation is set to 0 with probability p. If set to 0, then all downstream units receive no input from the unit, and the backpropagation algorithm will

not assign any portion of the calculated error to any of the ingoing or outgoing connections from the unit, negating learning for the associated weights on that trial and leaving those weights unchanged. The effect of dropout is a form of bagging for RNNs.

L2 weight decay ensures that only meaningful weights stay large and that weights are less biased by the specific training set by providing a proportionally increasing penalty to weights of large magnitude (Krogh and Hertz, 1992a). At each time step in a training iteration, weights are decreased by $d\mathbf{w}(t)$, where d is the weight decay rate, and $\mathbf{w}(t)$ is the weight at time step t.

Input resampling helps decrease the variance of the model by artificially increasing the size of the training set. We have defined input resampling to mean that at each time step in a training trial, with probability p, we substitute the remainder of the training trial time steps with another random same-target trial, matching electrodes and time steps. For example, while training on a trial in which the macaque successfully reached target B, we randomly draw a number below p on time step 7. Then, we choose a random trial in which the macaque also successfully reached target B and substitute the data for time steps 7 through 10 from that trial into the current training trial's time steps 7 through 10. Now we continue training as before, and at each successive time step we again determine whether to substitute a random same-target trial. We found an input resampling rate of 0.04 to be suitable for the generalized version of the RNN.

Probabilistic targets relax the assumption that the monkey always desires to saccade to the correct target for the entire 750 ms duration of the delay period. In a normal softmax output, the desired response is set to 1, and all other responses are set to

0 when the error gradient is calculated. With probabilistic targets, the correct target probability is set below 1, representing a Bayesian prior on the "correct" response, and the error gradient would therefore be smaller for incorrect responses. We found 0.87 to be a suitable value for the generalized version of the RNN. As far as we know, this generalization technique is novel.

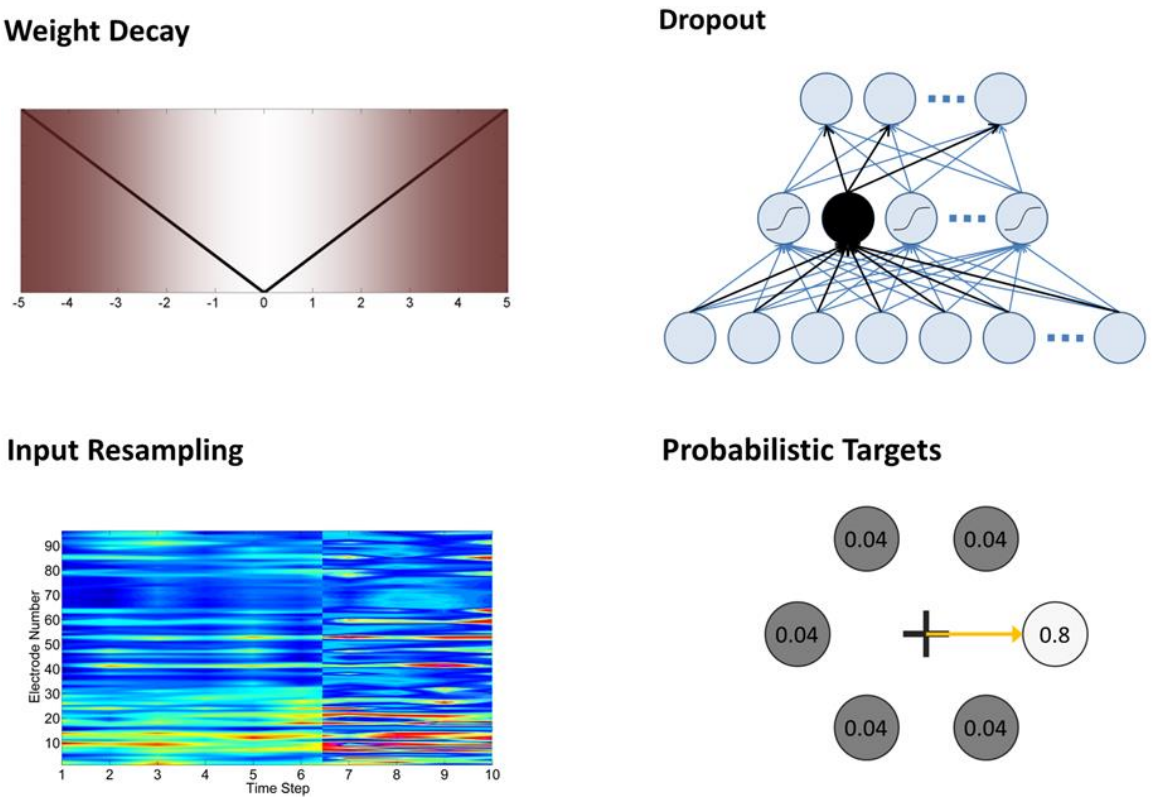

Figure 64: Generalization techniques used in RNN creation

Top left: L2 weight decay regularizes weight parameters to prevent many weights from growing too large. Top right: dropout simulates the training of $\mathbf{2^n}$ networks, approximating the bagging of many RNN architectures. Bottom left: input resampling effectively generates new training data, reducing the bias of having a limited training data set. Bottom right:

probabilistic targets approximate uncertainty in subject intention, preventing incorrectly labeled training data from severely impacting performance.

Saccade Behavior

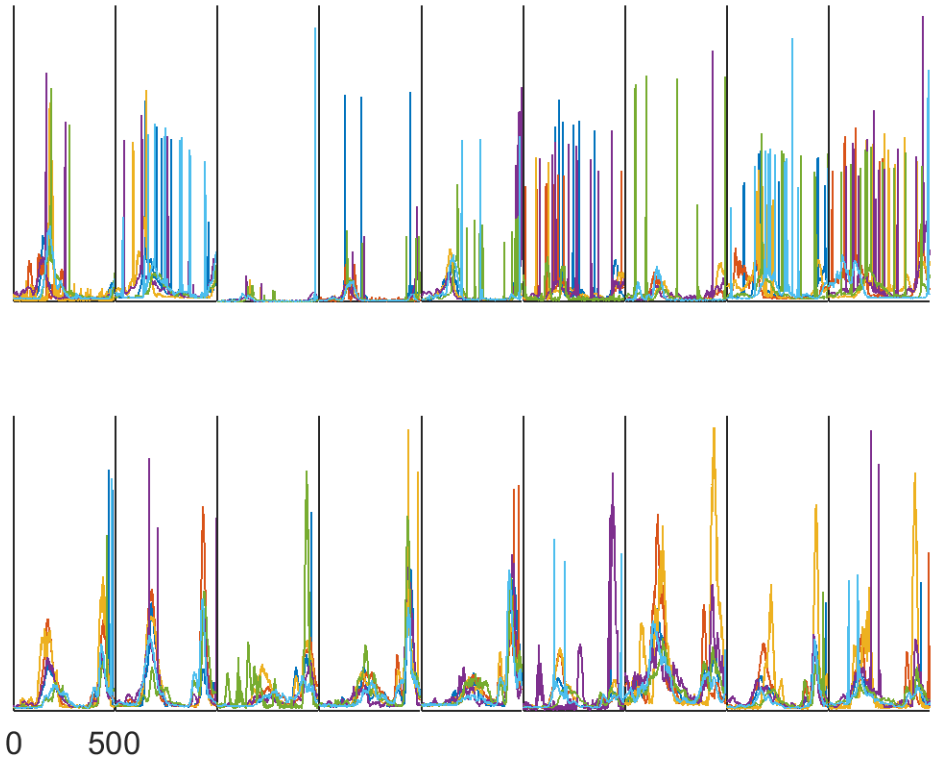

Figure 65: Most initial saccade activity complete by ~350 ms

Saccade velocity indicates that most movement to initial targets during both CORRECT and INCORRECT trials was completed around ~350 ms, and that there are often secondary targets after ~350 ms. Based on this

analysis, we chose to display only the first 350 ms after GO in Fig. COR/INC.

Reconstruction of 80-475 Hz Power

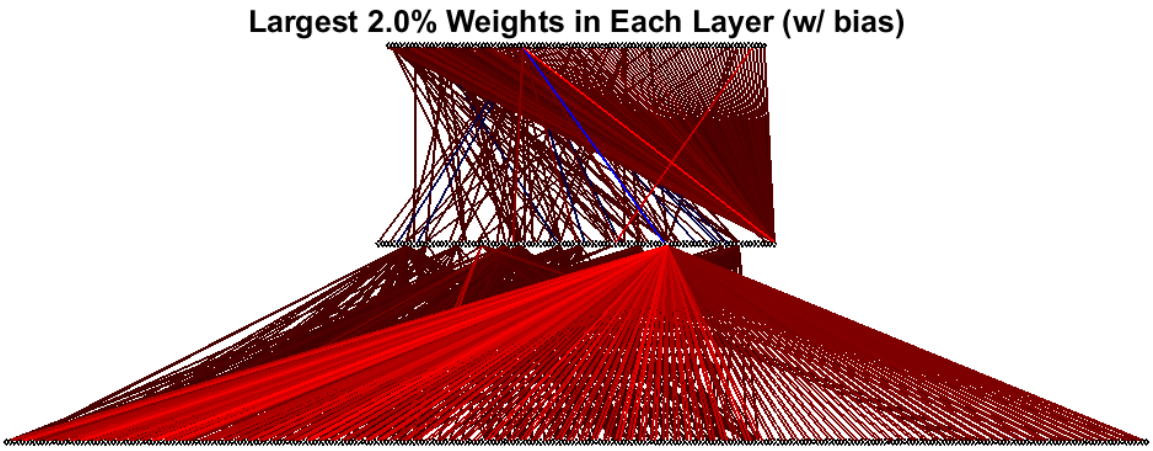

Figure 66: Reconstruction bias weights for monkey C are large

Largest 2% of weights (by absolute magnitude) in each layer for monkey C (session 20140304). This shows that the bias weights (which are included in this image) are generally large between the hidden layer and the output layer, indicating a power floor exists for 80-475 Hz. Note that most of the largest weights (bright red) are coming from the 1-2 Hz power inputs (first 1/3rd of inputs at left).

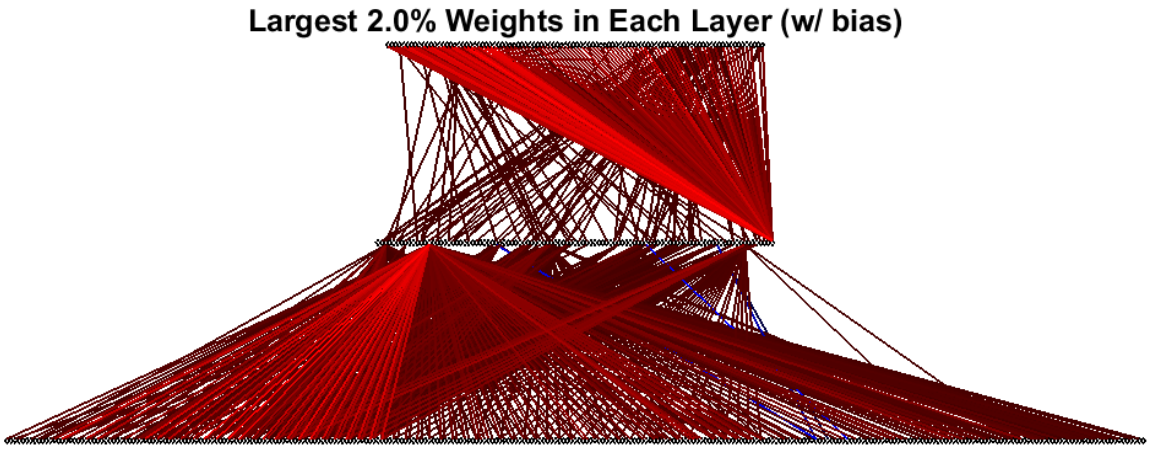

Figure 67: Reconstruction bias weights for monkey J are large

Largest 2% of weights (by absolute magnitude) in each layer for monkey J (session 20140304). This shows that the bias weights (which are included in this image) are generally large between the hidden layer and the output layer, indicating a power floor exists for 80-475 Hz. Note that most of the largest weights (bright red) are coming from the 1-2 Hz power inputs (first 1/3rd of inputs at left).

BIBLIOGRAPHY

Achtman N, Afshar A, Santhanam G, Yu BM, Ryu SI, Shenoy K V (2007) Free-paced high-performance brain-computer interfaces. Journal of Neural Engineering 4:336–347.

Ahrens KF, Kleinfeld D (2004) Current flow in vibrissa motor cortex can phase-lock with exploratory rhythmic whisking in rat. Journal of Neurophysiology 92:1700–1707.

Bai O, Lin P, Vorbach S, Floeter MK, Hattori N, Hallett M (2008) A high performance sensorimotor beta rhythm-based brain-computer interface associated with human natural motor behavior. Journal of Neural Engineering 5:24–35.

Baker MR, Baker SN (2002) The effect of diazepam on motor cortical oscillations and corticomuscular coherence studied in man. The Journal of Physiology 546:931–942.

Baker SN, Chiu M, Fetz EE (2006) Afferent encoding of central oscillations in the monkey arm. Journal of Neurophysiology 95:3904–3910.

Baker SN, Olivier E, Lemon RN (1997) Coherent oscillations in monkey motor cortex and hand muscle EMG show task-dependent modulation. The Journal of Physiology 501:225–241.

Baker SN, Pinches EM, Lemon RN (2003) Synchronization in monkey motor cortex during a precision grip task. II. effect of oscillatory activity on corticospinal output. Journal of Neurophysiology 89:1941–1953.

Bansal AK, Truccolo W, Vargas-Irwin CE, Donoghue JP (2012) Decoding 3D reach and grasp from hybrid signals in motor and premotor cortices: spikes, multiunit activity, and local field potentials. Journal of Neurophysiology 107:1337–1355.

Barton RR, Schruben LW (2001) Resampling methods for input modeling. In: Proceedings of the 2001 Winter Simulation Conference, pp.372–378.

Berg RW, Kleinfeld D (2003) Rhythmic whisking by rat: retraction as well as protraction of the vibrissae is under active muscular control. Journal of Neurophysiology 89:104–117.

Birch GE, Bozorgzadeh Z, Mason SG (2002) Initial on-line evaluations of the LF-ASD brain-computer interface with able-bodied and spinal-cord subjects using imagined voluntary motor potentials. IEEE Transactions on Neural Systems and Rehabilitation Engineering 10:219–224.

Bishop W, Chestek CC, Gilja V, Nuyujukian P, Foster JD, Ryu SI, Shenoy K V, Yu BM (2014) Self-recalibrating classifiers for intracortical brain-computer interfaces.

Journal of Neural Engineering 11:26001.

Boulay CB, Pieper F, Leavitt M, Martinez-Trujillo J, Sachs AJ (2016) Single-trial decoding of intended eye movement goals from lateral prefrontal cortex neural ensembles. Journal of Neurophysiology 115:486–499.

Boulay CB, Sarnacki WA, Wood F, McFarland DJ (2011) Trained modulation of sensorimotor rhythms can affect reaction time. Clinical Neurophysiology 122:1820–1826.

Brecht M, Schneider M, Sakmann B, Margrie TW (2004) Whisker movements evoked by stimulation of single pyramidal cells in rat motor cortex. Nature 427:704–710.

Brincat SL, Jia N, Salazar-Gomez AF, Panko M, Miller EK, Guenther FH (2013a) Which neural signals are optimal for brain-computer interface control? In: Proceedings of the Fifth International Brain-Computer Interface Meeting 2013, pp.104.

Brincat SL, Salazar-Gomez AF, Jia N, Panko M, Miller EK, Guenther FH (2013b) Development of an Intracortical Eye Movement-Based Brain-Computer Interface. In: Proceedings of the Fifth International Brain-Computer Interface Meeting 2013, pp.34.

Bruce CJ, Goldberg ME (1985) Primate frontal eye fields. I. Single neurons discharging before saccades. Journal of Neurophysiology 53:603–635.

Bryan MJ, Martin S a, Cheung W, Rao RPN (2013) Probabilistic co-adaptive brain-computer interfacing. Journal of Neural Engineering 10:66008.

Bundy DT, Pahwa M, Szrama N, Leuthardt EC (2016) Decoding three-dimensional reaching movements using electrocorticographic signals in humans. Journal of Neural Engineering 13:26021.

Buschman TJ, Denovellis EL, Diogo C, Bullock D, Miller EK (2012) Synchronous oscillatory neural ensembles for rules in the prefrontal cortex. Neuron 76:838–846.

Cameron IGM, Riddle JM, D'Esposito M (2015) Dissociable roles of dorsolateral prefrontal cortex and frontal eye fields during saccadic eye movements. Frontiers in Human Neuroscience 9:1–14.

Canolty RT, Ganguly K, Carmena JM (2012) Task-dependent changes in cross-level coupling between single neurons and oscillatory activity in multiscale networks. PLoS Computational Biology 8:e1002809.

Cao L, Hu Y (2016) Beta Rebound in Visuomotor Adaptation: Still the Status Quo? The Journal of Neuroscience 36:6365–6367.

Cao Y, Roy S, Sachdev RNS, Heck DH (2012) Dynamic correlation between whisking and breathing rhythms in mice. The Journal of Neuroscience 32:1653–1659.

Carmena JM, Lebedev MA, Crist RE, O'Doherty JE, Santucci DM, Dimitrov DF, Patil PG, Henriquez CS, Nicolelis MAL (2003) Learning to control a brain-machine interface for reaching and grasping by primates. PLoS Biology 1:193–208.

Cassidy M, Mazzone P, Oliviero A, Insola A, Tonali P, Di Lazzaro V, Brown P (2002) Movement-related changes in synchronization in the human basal ganglia. Brain 125:1235–1246.

Cassim F, Monaca C, Szurhaj W, Bourriez J, Defebvre L, Derambure P, Guieu J-D (2001) Does post-movement beta synchronization reflect an idling motor cortex? NeuroReport 12:3859–3863.

Castro-Alamancos MA (2006) Vibrissa myoclonus (rhythmic retractions) driven by resonance of excitatory networks in motor cortex. Journal of Neurophysiology 96:1691–1698.

Cerf M, Thiruvengadam N, Mormann F, Kraskov A, Quiroga RQ, Koch C, Fried I (2010) On-line, voluntary control of human temporal lobe neurons. Nature 467:1104–1108.

Chao ZC, Nagasaka Y, Fujii N (2010) Long-term asynchronous decoding of arm motion using electrocorticographic signals in monkeys. Frontiers in Neuroengineering 3:3.

Chapin JK, Moxon KA, Markowitz RS, Nicolelis MAL (1999) Real-time control of a robot arm using simultaneously recorded neurons in the motor cortex. Nature Neuroscience 2:664–670.

Chase SM, Kass RE, Schwartz AB (2012) Behavioral and neural correlates of visuomotor adaptation observed through a brain-computer interface in primary motor cortex. Journal of Neurophysiology 108:624–644.

Chase SM, Schwartz AB, Kass RE (2009) Bias, optimal linear estimation, and the differences between open-loop simulation and closed-loop performance of spiking-based brain-computer interface algorithms. Neural Networks 22:1203–1213.

Chestek CA, Gilja V, Nuyujukian P, Foster JD, Fan JM, Kaufman MT, Churchland MM, Rivera-Alvidrez Z, Cunningham JP, Ryu SI, Shenoy K V (2011) Long-term stability of neural prosthetic control signals from silicon cortical arrays in rhesus macaque motor cortex. Journal of Neural Engineering 8:45005.

Clancy KB, Koralek AC, Costa RM, Feldman DE, Carmena JM (2014) Volitional modulation of optically recorded calcium signals during neuroprosthetic learning. Nature Neuroscience 17:807–809.

Collinger JL, Wodlinger B, Downey JE, Wang W, Tyler-Kabara EC, Weber DJ, McMorland AJ, Velliste M, Boninger ML, Schwartz AB (2013) High-performance neuroprosthetic control by an individual with tetraplegia. The Lancet 381:557–564.

Cunningham JP, Nuyujukian P, Gilja V, Chestek C a, Ryu SI, Shenoy K V (2011) A closed-loop human simulator for investigating the role of feedback control in brain-machine interfaces. Journal of Neurophysiology 105:1932–1949.

Daly JJ, Wolpaw JR (2008) Brain-computer interfaces in neurological rehabilitation. The Lancet Neurology 7:1032–1043.

Dangi S, Gowda S, Moorman HG, Orsborn AL, So K, Shanechi MM, Carmena JM (2014) Continuous closed-loop decoder adaptation with a recursive maximum likelihood algorithm allows for rapid performance acquisition in brain-machine interfaces. Neural Computation 26:1811–1839.

Dangi S, Orsborn AL, Moorman HG, Carmena JM (2013) Design and analysis of closed-loop decoder adaptation algorithms for brain-machine interfaces. Neural Computation 25:1693–1731.

Danziger Z, Fishbach A, Mussa-Ivaldi FA (2009) Learning algorithms for human-machine interfaces. IEEE Transactions on Biomedical Engineering 56:1502–1511.

DaSalla CS, Kambara H, Sato M, Koike Y (2009) Single-trial classification of vowel speech imagery using common spatial patterns. Neural Networks 22:1334–1339.

Dias EC, Bruce CJ (1994) Physiological correlate of fixation disengagement in the primate's frontal eye field. Journal of Neurophysiology 72:2532–2537.

DiGiovanna J, Mahmoudi B, Fortes J, Principe JC, Sanchez JC (2009) Coadaptive Brain-Machine Interface via Reinforcement Learning. IEEE Transactions on Bio-Medical Engineering 56:54–64.

Donoghue JP, Sanes JN, Hatsopoulos NG, Gaál G (1998) Neural Discharge and Local Field Potential Oscillations in Primate Motor Cortex During Voluntary Movements. Journal of Neurophysiology 79:159–173.

Dufort RH, Abrahamson DS (1966) Water intake of the rat as a function of duration of water deprivation. Psychological Reports 18:148–150.

van Ede F, Maris E (2013) Somatosensory Demands Modulate Muscular Beta Oscillations, Independent of Motor Demands. Journal of Neuroscience 33:10849–10857.

Elman JL (1990) Finding structure in time. Cognitive Science 14:179–211.

Engel AK, Fries P (2010) Beta-band oscillations--signalling the status quo? Current Opinion in Neurobiology 20:156–165.

Engelhard B, Ozeri N, Israel Z, Bergman H, Vaadia E (2013) Inducing Gamma Oscillations and Precise Spike Synchrony by Operant Conditioning via Brain-Machine Interface. Neuron 77:361–375.

Erlich JC, Bialek M, Brody CD (2011) A cortical substrate for memory-guided orienting in the rat. Neuron 72:330–343.

Faisal AA, Selen LPJ, Wolpert DM (2008) Noise in the nervous system. Nature Reviews Neuroscience 9:292–303.

Fan JM, Nuyujukian P, Kao JC, Chestek CA, Ryu SI, Shenoy K V (2014) Intention estimation in brain–machine interfaces. Journal of Neural Engineering 11:16004.

Farkas T, Kis Z, Toldi J, Wolff JR (1999) Activation of the primary motor cortex by somatosensory stimulation in adult rats is mediated mainly by associational connections from the somatosensory cortex. Neuroscience 90:353–361.

Fatourechi M, Ward RK, Birch GE (2008) A self-paced brain-computer interface system with a low false positive rate. Journal of Neural Engineering 5:9–23.

Feingold J, Gibson DJ, DePasquale B, Graybiel AM (2015) Bursts of beta oscillation differentiate postperformance activity in the striatum and motor cortex of monkeys performing movement tasks. Proceedings of the National Academy of Sciences 112:13687–13692.

Fernandez-Vargas J, Pfaff HU, Rodríguez FB, Varona P (2013) Assisted closed-loop optimization of SSVEP-BCI efficiency. Frontiers in Neural Circuits 7:27.

Fetz EE (1969) Operant Conditioning of Cortical Unit Activity. Science 163:955–958.

Fetz EE, Baker MA (1973) Operantly conditioned patterns of precentral unit activity and correlated responses in adjacent cells and contralateral muscles. Journal of Neurophysiology 36:179–204.

Flint RD, Ethier C, Oby ER, Miller LE, Slutzky MW (2012) Local field potentials allow accurate decoding of muscle activity. Journal of Neurophysiology 108:18–24.

Flint RD, Wright Z a, Scheid MR, Slutzky MW (2013) Long term, stable brain machine interface performance using local field potentials and multiunit spikes. Journal of Neural Engineering 10:56005.

Fontolan L, Morillon B, Liegeois-Chauvel C, Giraud A-L (2014) The contribution of frequency-specific activity to hierarchical information processing in the human

auditory cortex. Nature Communications 5:4694.

Fraser GW, Chase SM, Whitford A, Schwartz AB (2009) Control of a brain-computer interface without spike sorting. Journal of Neural Engineering 6:55004.

Friston KJ, Bastos AM, Pinotsis D, Litvak V (2015) LFP and oscillations-what do they tell us? Current Opinion in Neurobiology 31:1–6.

Fujioka T, Trainor LJ, Large EW, Ross B (2012) Internalized timing of isochronous sounds is represented in neuromagnetic β oscillations. The Journal of Neuroscience 32:1791–1802.

Funahashi S, Bruce CJ, Goldman-Rakic PS (1989) Mnemonic Coding of Visual Space in the Monkey's Dorsolateral Prefrontal Cortex. Journal of Neurophysiology 61:331–349.

Funahashi S, Bruce CJ, Goldman-Rakic PS (1993) Dorsolateral prefrontal lesions and oculomotor delayed-response performance: evidence for mnemonic "scotomas." The Journal of Neuroscience 13:1479–1497.

Gage GJ, Ludwig KA, Otto KJ, Ionides EL, Kipke DR (2005) Naive coadaptive cortical control. Journal of Neural Engineering 2:52–63.

Ganguly K, Carmena JM (2009) Emergence of a stable cortical map for neuroprosthetic control. PLoS Biology 7:e1000153.

Ganguly K, Dimitrov DF, Wallis JD, Carmena JM (2011) Reversible large-scale modification of cortical networks during neuroprosthetic control. Nature Neuroscience 14:662–667.

Ganguly K, Kleinfeld D (2004) Goal-directed whisking increases phase-locking between vibrissa movement and electrical activity in primary sensory cortex in rat. Proceedings of the National Academy of Sciences of the United States of America 101:12348–12353.

Gao R (2016) Interpreting the Electrophysiological Power Spectrum. Journal of Neurophysiology 115:628–630.

Gilja V, Nuyujukian P, Chestek CA, Cunningham JP, Yu BM, Fan JM, Churchland MM, Kaufman MT, Kao JC, Ryu SI, Shenoy K V (2012) A high-performance neural prosthesis enabled by control algorithm design. Nature Neuroscience 15:1752–1757.

Graf ABA, Andersen RA (2014) Brain–machine interface for eye movements. Proceedings of the National Academy of Sciences 111:17630–17635.

Grosse-Wentrup M, Schölkopf B (2014) A brain-computer interface based on self-

regulation of gamma-oscillations in the superior parietal cortex. Journal of Neural Engineering 11:56015.

Guclu U, van Gerven MAJ (2015) Deep Neural Networks Reveal a Gradient in the Complexity of Neural Representations across the Ventral Stream. Journal of Neuroscience 35:10005–10014.

Guenther FH, Brumberg JS, Wright EJ, Nieto-Castanon A, Tourville JA, Panko M, Law R, Siebert SA, Bartels JL, Andreasen DS, Ehirim P, Mao H, Kennedy PR (2009) A wireless brain-machine interface for real-time speech synthesis. PLoS ONE 4:e8218.

Gürel T, Mehring C (2012) Unsupervised adaptation of brain-machine interface decoders. Frontiers in Neuroscience 6:164.

Hamada Y, Miyashita E, Tanaka H (1999) Gamma-band oscillations in the "barrel cortex" precede rat's exploratory whisking. Neuroscience 88:667–671.

Harris KD, Henze DA, Csicsvari J, Hirase H, Buzsaki G (2000) Accuracy of tetrode spike separation as determined by simultaneous intracellular and extracellular measurements. Journal of Neurophysiology 84:401–414.

Hill DN, Curtis JC, Moore JD, Kleinfeld D (2011) Primary Motor Cortex Reports Efferent Control of Vibrissa Motion on Multiple Timescales. Neuron 72:344–356.

Hochberg LR, Bacher D, Jarosiewicz B, Masse NY, Simeral JD, Vogel J, Haddadin S, Liu J, Cash SS, Smagt P Van Der, Donoghue JP (2012) Reach and grasp by people with tetraplegia using a neurally controlled robotic arm. Nature 485:372–375.

Holdefer RN, Miller LE (2002) Primary motor cortical neurons encode functional muscle synergies. Experimental Brain Research 146:233–243.

Hwang EJ, Andersen RA (2009) Brain control of movement execution onset using local field potentials in posterior parietal cortex. The Journal of Neuroscience 29:14363–14370.

Jarosiewicz B, Chase SM, Fraser GW, Velliste M, Kass RE, Schwartz AB (2008) Functional network reorganization during learning in a brain-computer interface paradigm. Proceedings of the National Academy of Sciences of the United States of America 105:19486–19491.

Jin XL, Guo H, Mao C, Atkins N, Wang H, Avasthi PP, Tu YT, Li Y (2000) Emx1-specific expression of foreign genes using "knock-in" approach. Biochemical and Biophysical Research Communications 270:978–982.

Jones SR (2016) When brain rhythms aren't "rhythmic": implication for their

mechanisms and meaning. Current Opinion in Neurobiology 40:72–80.

Jones SR, Pritchett DL, Sikora M a, Stufflebeam SM, Hämäläinen M, Moore CI (2009) Quantitative analysis and biophysically realistic neural modeling of the MEG mu rhythm: rhythmogenesis and modulation of sensory-evoked responses. Journal of Neurophysiology 102:3554–3572.

Jordan MI (1986) Serial order: A parallel distributed processing approach. University of California, San Diego Institute for Cognitive Science:8604.

Joundi RA, Jenkinson N, Brittain J-S, Aziz TZ, Brown P (2012) Driving oscillatory activity in the human cortex enhances motor performance. Current Biology 22:403–407.

Kao JC, Nuyujukian P, Ryu SI, Churchland MM, Cunningham JP, Shenoy K V (2015) Single-trial dynamics of motor cortex and their applications to brain-machine interfaces. Nature Communications 6:1–12.

Kaufman MT, Churchland MM, Ryu SI, Shenoy K V (2014) Cortical activity in the null space: permitting preparation without movement. Nature Neuroscience 17:440–448.

Khanna P, Carmena JM (2015) Effects of volitional modulation of beta oscillations on reaching tasks. In: Neuroscience Meeting Planner, pp.427.14.

Kilavik BE, Ponce-Alvarez A, Trachel R, Confais J, Takerkart S, Riehle A (2012) Context-related frequency modulations of macaque motor cortical LFP beta oscillations. Cerebral Cortex 22:2148–2159.

Kilner JM, Fisher RJ, Lemon RN (2004) Coupling of oscillatory activity between muscles is strikingly reduced in a deafferented subject compared with normal controls. Journal of Neurophysiology 92:790–796.

Kim S-P, Simeral JD, Hochberg LR, Donoghue JP, Black MJ (2008) Neural control of computer cursor velocity by decoding motor cortical spiking activity in humans with tetraplegia. Journal of Neural Engineering 5:455–476.

Kleinfeld D, Deschênes M, Wang F, Moore JD (2014) More than a rhythm of life: breathing as a binder of orofacial sensation. Nature Neuroscience 17:647–651.

Kobayashi S, Schultz W, Sakagami M (2010) Operant conditioning of primate prefrontal neurons. Journal of Neurophysiology 103:1843–1855.

Koralek AC, Jin X, Long Ii JD, Costa RM, Carmena JM (2012) Corticostriatal plasticity is necessary for learning intentional neuroprosthetic skills. Nature 483:331–335.

Koyama S, Chase SM, Whitford AS, Velliste M, Schwartz AB, Kass RE (2010)

Comparison of brain-computer interface decoding algorithms in open-loop and closed-loop control. Journal of Computational Neuroscience 29:73–87.

Krogh A, Hertz JA (1992a) A Simple Weight Decay Can Improve Generalization. Advances in Neural Information Processing Systems 4:950–957.

Krogh A, Hertz JA (1992b) Generalization in a linear perceptron in the presence of noise. Journal of Physics A: Mathematical and General 25:1135–1147.

Kropotov JD, Grin-Yatsenko VA, Ponomarev VA, Chutko LS, Yakovenko EA, Nikishena IS (2005) ERPs correlates of EEG relative beta training in ADHD children. International Journal of Psychophysiology 55:23–34.

Kübler A, Birbaumer NP (2008) Brain-computer interfaces and communication in paralysis: extinction of goal directed thinking in completely paralysed patients? Clinical Neurophysiology 119:2658–2666.

Kühn AA, Williams D, Kupsch A, Limousin P, Hariz M, Schneider G-H, Yarrow K, Brown P (2004) Event-related beta desynchronization in human subthalamic nucleus correlates with motor performance. Brain 127:735–746.

Lebedev MA, Carmena JM, O'Doherty JE, Zacksenhouse M, Henriquez CS, Principe JC, Nicolelis MAL (2005) Cortical ensemble adaptation to represent velocity of an artificial actuator controlled by a brain-machine interface. The Journal of Neuroscience 25:4681–4693.

Leeb R, Friedman D, Müller-Putz GR, Scherer R, Slater M, Pfurtscheller G (2007) Self-paced (asynchronous) BCI control of a wheelchair in virtual environments: a case study with a tetraplegic. Computational Intelligence and Neuroscience 2007:79642.

Leventhal DK, Gage GJ, Schmidt R, Pettibone JR, Case AC, Berke JD (2012) Basal ganglia beta oscillations accompany cue utilization. Neuron 73:523–536.

Li Z, O'Doherty JE, Lebedev MA, Nicolelis MAL (2011) Adaptive decoding for brain-machine interfaces through Bayesian parameter updates. Neural Computation 23:3162–3204.

Liu Y, Denton JM, Nelson RJ (2005) Neuronal activity in primary motor cortex differs when monkeys perform somatosensory and visually guided wrist movements. Experimental Brain Research 167:571–586.

Logothetis NK (2003) The underpinnings of the BOLD functional magnetic resonance imaging signal. The Journal of Neuroscience 23:3963–3971.

Ludwig KA, Miriani RM, Langhals NB, Marzullo TC, Kipke DR (2011) Use of a

Bayesian maximum-likelihood classifier to generate training data for brain–machine interfaces. Journal of Neural Engineering 8:49801.

Lundqvist M, Rose J, Herman P, Brincat SL, Buschman TJ, Miller EK (2016) Gamma and Beta Bursts Underlie Working Memory. Neuron 90:152–164.

MacKay W, Mendonça A (1995) Field potential oscillatory bursts in parietal cortex before and during reach. Brain Research 704:167–174.

Madisen L et al. (2012) A toolbox of Cre-dependent optogenetic transgenic mice for light-induced activation and silencing. Nature Neuroscience 15:793–802.

Manning JR, Jacobs J, Fried I, Kahana MJ (2009) Broadband shifts in local field potential power spectra are correlated with single-neuron spiking in humans. The Journal of Neuroscience 29:13613–13620.

Mao T, Kusefoglu D, Hooks BM, Huber D, Petreanu L, Svoboda K (2011) Long-range neuronal circuits underlying the interaction between sensory and motor cortex. Neuron 72:111–123.

Marsh BT, Tarigoppula VSA, Chen C, Francis JT (2015) Toward an Autonomous Brain Machine Interface: Integrating Sensorimotor Reward Modulation and Reinforcement Learning. Journal of Neuroscience 35:7374–7387.

Mason SG, Birch GE (2000) A brain-controlled switch for asynchronous control applications. IEEE Transactions on Biomedical Engineering 47:1297–1307.

Matyas F, Sreenivasan V, Marbach F, Wacongne C, Barsy B, Mateo C, Aronoff R, Petersen CCH (2010) Motor control by sensory cortex. Science 330:1240–1243.

McCreadie K a, Coyle DH, Prasad G (2013) Sensorimotor learning with stereo auditory feedback for a brain-computer interface. Medical & Biological Engineering & Computing 51:285–293.

McFarland DJ, Sarnacki WA, Wolpaw JR (2015) Effects of training pre-movement sensorimotor rhythms on behavioral performance. Journal of Neural Engineering 12:66021.

Merel J, Pianto DM, Cunningham JP, Paninski L (2015) Encoder-Decoder Optimization for Brain-Computer Interfaces. PLOS Computational Biology 11:e1004288.

Michalareas G, Vezoli J, van Pelt S, Schoffelen J-M, Kennedy H, Fries P (2016) Alpha-Beta and Gamma Rhythms Subserve Feedback and Feedforward Influences among Human Visual Cortical Areas. Neuron 89:384–397.

Miller KJ (2010) Broadband spectral change: evidence for a macroscale correlate of

population firing rate? The Journal of Neuroscience 30:6477–6479.

Miller KJ, Schalk G, Fetz EE, den Nijs M, Ojemann JG, Rao RPN (2010) Cortical activity during motor execution, motor imagery, and imagery-based online feedback. Proceedings of the National Academy of Sciences of the United States of America 107:4430–4435.

Miller KJ, Sorensen LB, Ojemann JG, den Nijs M (2009a) Power-law scaling in the brain surface electric potential. PLoS Computational Biology 5:e1000609.

Miller KJ, Zanos S, Fetz EE, den Nijs M, Ojemann JG (2009b) Decoupling the cortical power spectrum reveals real-time representation of individual finger movements in humans. The Journal of Neuroscience 29:3132–3137.

Mohler CW, Goldberg ME, Wurtz RH (1973) Visual receptive fields of frontal eye field neurons. Brain Research 61:385–389.

Moritz CT, Fetz EE (2011) Volitional control of single cortical neurons in a brain-machine interface. Journal of Neural Engineering 8:25017.

Murthy V, Fetz EE (1996) Oscillatory activity in sensorimotor cortex of awake monkeys: synchronization of local field potentials and relation to behavior. Journal of Neurophysiology 76:3949–3967.

Murthy VN, Fetz EE (1992) Coherent 25- to 35-Hz oscillations in the sensorimotor cortex of awake behaving monkeys. Proceedings of the National Academy of Sciences of the United States of America 89:5670–5674.

Musallam S, Corneil B, Greger B, Scherberger H, Andersen RA (2004) Cognitive control signals for neural prosthetics. Science 305:258–262.

Naros G, Gharabaghi A (2015) Reinforcement learning of self-regulated β-oscillations for motor restoration in chronic stroke. Frontiers in Human Neuroscience 9:391.

do Nascimento OF, Nielsen KD, Voigt M (2006) Movement-related parameters modulate cortical activity during imaginary isometric plantar-flexions. Experimental Brain Research 171:78–90.

O'Leary JG, Hatsopoulos NG (2006) Early visuomotor representations revealed from evoked local field potentials in motor and premotor cortical areas. Journal of Neurophysiology 96:1492–1506.

Ohmae S, Takahashi T, Lu X, Nishimori Y, Kodaka Y, Takashima I, Kitazawa S (2015) Decoding the timing and target locations of saccadic eye movements from neuronal activity in macaque oculomotor areas. Journal of Neural Engineering 12:36014.

Omlor W, Patino L, Hepp-Reymond M-C, Kristeva R (2007) Gamma-range corticomuscular coherence during dynamic force output. NeuroImage 34:1191–1198.

Orsborn AL, Dangi S, Moorman HG, Carmena JM (2011) Exploring time-scales of closed-loop decoder adaptation in brain-machine interfaces. In: 33rd Annual International Conference of the IEEE EMBS, pp.5436–5439.

Orsborn AL, Dangi S, Moorman HG, Carmena JM (2012) Closed-loop decoder adaptation on intermediate time-scales facilitates rapid BMI performance improvements independent of decoder initialization conditions. IEEE Transactions on Neural Systems and Rehabilitation Engineering 20:468–477.

Orsborn AL, Moorman HG, Overduin SA, Shanechi MM, Dimitrov DF, Carmena JM (2014) Closed-Loop Decoder Adaptation Shapes Neural Plasticity for Skillful Neuroprosthetic Control. Neuron 82:1380–1393.

Panzeri S, Safaai H, De Feo V, Vato A (2016) Implications of the Dependence of Neuronal Activity on Neural Network States for the Design of Brain-Machine Interfaces. Frontiers in Neuroscience 10:165.

Paré M, Munoz DP (1996) Saccadic reaction time in the monkey: advanced preparation of oculomotor programs is primarily responsible for express saccade occurrence. Journal of Neurophysiology 76:3666–3681.

Parikh H, Marzullo TC, Kipke DR (2009) Lower layers in the motor cortex are more effective targets for penetrating microelectrodes in cortical prostheses. Journal of Neural Engineering 6:26004.

Parkkonen E, Laaksonen K, Piitulainen H, Parkkonen L, Forss N (2015) Modulation of the ~20-Hz motor-cortex rhythm to passive movement and tactile stimulation. Brain and Behavior 5:e00328.

Paxinos G, Franklin KBJ (2001) The mouse brain in stereotaxic coordinates (2nd edition).

Perge JA, Homer ML, Malik WQ, Cash S, Eskandar E, Friehs G, Donoghue JP, Hochberg LR (2013) Intra-day signal instabilities affect decoding performance in an intracortical neural interface system. Journal of Neural Engineering 10:36004.

Pesaran B, Pezaris JS, Sahani M, Mitra PP, Andersen RA (2002) Temporal structure in neuronal activity during working memory in macaque parietal cortex. Nature Neuroscience 5:805–811.

Pfurtscheller G, Lopes da Silva F (1999) Event-related EEG/MEG synchronization and

desynchronization: basic principles. Clinical Neurophysiology 110:1842–1857.

Pfurtscheller G, Neuper C, Brunner C, Lopes da Silva F (2005) Beta rebound after different types of motor imagery in man. Neuroscience Letters 378:156–159.

Pfurtscheller G, Solis-Escalante T (2009) Could the beta rebound in the EEG be suitable to realize a "brain switch"? Clinical Neurophysiology 120:24–29.

Pfurtscheller G, Solis-Escalante T, Ortner R, Linortner P, Müller-Putz GR (2010) Self-paced operation of an SSVEP-Based orthosis with and without an imagery-based "brain switch:" a feasibility study towards a hybrid BCI. IEEE Transactions on Neural Systems and Rehabilitation Engineering 18:409–414.

Philippens IHCHM, Vanwersch RAP (2010) Neurofeedback training on sensorimotor rhythm in marmoset monkeys. Neuroreport 21:328–332.

Pichiorri F, De Vico Fallani F, Cincotti F, Babiloni F, Molinari M, Kleih SC, Neuper C, Kübler A, Mattia D (2011) Sensorimotor rhythm-based brain-computer interface training: the impact on motor cortical responsiveness. Journal of Neural Engineering 8:25020.

Prasad A, Xue Q-S, Sankar V, Nishida T, Shaw G, Streit WJ, Sanchez JC (2012) Comprehensive characterization and failure modes of tungsten microwire arrays in chronic neural implants. Journal of Neural Engineering 9:56015.

Pronichev I V, Lenkov DN (1998) Functional mapping of the motor cortex of the white mouse by a microstimulation method. Neuroscience and Behavioral Physiology 28:80–85.

Pruszynski JA, Kurtzer I, Scott SH (2011) The long-latency reflex is composed of at least two functionally independent processes. Journal of Neurophysiology 106:449–459.

Quiroga RQ, Nadasdy Z, Ben-Shaul Y (2004) Unsupervised spike detection and sorting with wavelets and superparamagnetic clustering. Neural Computation 16:1661–1687.

Ramot M, Grossman S, Friedman D, Malach R (2016) Covert neurofeedback without awareness shapes cortical network spontaneous connectivity. Proceedings of the National Academy of Sciences:In press.

Ray S, Maunsell JHR (2011) Different origins of gamma rhythm and high-gamma activity in macaque visual cortex. PLoS Biology 9:e1000610.

Rouse AG, Williams JJ, Wheeler JJ, Moran DW (2013) Cortical adaptation to a chronic micro-electrocorticographic brain computer interface. The Journal of Neuroscience

33:1326–1330.

Rubino D, Robbins KA, Hatsopoulos NG (2006) Propagating waves mediate information transfer in the motor cortex. Nature Neuroscience 9:1549–1557.

Ryu SI, Shenoy K V (2009) Human cortical prostheses: lost in translation? Neurosurgical Focus 27:E5.

Sacchet MD, Laplante RA, Wan Q, Pritchett DL, Lee AKC, Hämäläinen M, Moore CI, Kerr CE, Jones SR (2015) Attention Drives Synchronization of Alpha and Beta Rhythms between Right Inferior Frontal and Primary Sensory Neocortex. The Journal of Neuroscience 35:2074–2082.

Sadtler PT, Quick KM, Golub MD, Chase SM, Ryu SI, Tyler-Kabara EC, Yu BM, Batista AP (2014) Neural constraints on learning. Nature 512:423–426.

Sakurai Y, Takahashi S (2013) Conditioned enhancement of firing rates and synchrony of hippocampal neurons and firing rates of motor cortical neurons in rats. The European Journal of Neuroscience 37:623–639.

Sanchez JC, Mahmoudi B, DiGiovanna J, Principe JC (2009) Exploiting co-adaptation for the design of symbiotic neuroprosthetic assistants. Neural Networks 22:305–315.

Sanes JN, Donoghue JP (1993) Oscillations in local field potentials of the primate motor cortex during voluntary movement. Proceedings of the National Academy of Sciences of the United States of America 90:4470–4474.

Schafer RJ, Moore T (2011) Selective Attention from Voluntary Control of Neurons in Prefrontal Cortex. Science 332:1568–1572.

Schalk G, Kubánek J, Miller KJ, Anderson NR, Leuthardt EC, Ojemann JG, Limbrick D, Moran DW, Gerhardt L a, Wood F (2007) Decoding two-dimensional movement trajectories using electrocorticographic signals in humans. Journal of Neural Engineering 4:264–275.

Schall JD (1991) Neuronal activity related to visually guided saccadic eye movements in the supplementary motor area of rhesus monkeys. Journal of Neurophysiology 66:530–558.

Scherberger H, Jarvis MR, Andersen RA (2005) Cortical local field potential encodes movement intentions in the posterior parietal cortex. Neuron 46:347–354.

Schoffelen J-M, Oostenveld R, Fries P (2005) Neuronal coherence as a mechanism of effective corticospinal interaction. Science 308:111–113.

Schoffelen J-M, Poort J, Oostenveld R, Fries P (2011) Selective movement preparation is

subserved by selective increases in corticomuscular gamma-band coherence. The Journal of Neuroscience 31:6750–6758.

Schroeder JB, Ritt JT (2013) Extraction of intended palpation times from facial EMGs in a mouse model of active sensing. In: 35th Annual International Conference of the IEEE Engineering in Medicine and Biology Society, pp.2016–2019.

Shanechi MM, Hu RC, Powers M, Wornell GW, Brown EN, Williams ZM (2012) Neural population partitioning and a concurrent brain-machine interface for sequential motor function. Nature Neuroscience 15:1715–1722.

Shaw F, Chew J (2003) Dynamic changes of gamma activities of somatic cortical evoked potentials during wake – sleep states in rats. Brain Research 983:152–161.

Shenoy K V, Carmena JM (2014) Combining Decoder Design and Neural Adaptation in Brain-Machine Interfaces. Neuron 84:665–680.

Shenoy K V, Meeker D, Cao S, Kureshi SA, Pesaran B, Buneo CA, Batista AP, Mitra PP, Burdick JW, Andersen RA (2003) Neural prosthetic control signals from plan activity. Neuroreport 14:591–596.

Sherman MA, Lee S, Law R, Haegens S, Thorn CA, Hämäläinen MS, Moore CI, Jones SR (2016) Neural mechanisms of transient neocortical beta rhythms: Converging evidence from humans, computational modeling, monkeys, and mice. Proceedings of the National Academy of Sciences:in press.

Shpigelman L, Lalazar H, Vaadia E (2008) Kernel-ARMA for hand tracking and brain-machine interfacing during 3D motor control. Advances in neural …:1–8.

So K, Dangi S, Orsborn AL, Gastpar MC, Carmena JM (2014) Subject-specific modulation of local field potential spectral power during brain-machine interface control in primates. Journal of Neural Engineering 11:26002.

Solis-Escalante T, Müller-Putz G, Brunner C, Kaiser V, Pfurtscheller G (2010) Analysis of sensorimotor rhythms for the implementation of a brain switch for healthy subjects. Biomedical Signal Processing and Control 5:15–20.

Srivastava N, Hinton G, Krizhevsky A, Sutskever I, Salakhutdinov R (2014) Dropout: A Simple Way to Prevent Neural Networks from Overfitting. Journal of Machine Learning Research 15:1929–1958.

Stančák A, Riml A, Pfurtscheller G (1997) The effects of external load on movement-related changes of the sensorimotor EEG rhythms. Electroencephalography and Clinical Neurophysiology 102:495–504.

Stark E, Abeles M (2007) Predicting movement from multiunit activity. The Journal of Neuroscience 27:8387–8394.

Stellar E, Hill JH (1952) The rat's rate of drinking as a function of water deprivation. Journal of Comparative and Physiological Psychology 45:96–102.

Suminski AJ, Tkach DC, Fagg AH, Hatsopoulos NG (2010) Incorporating feedback from multiple sensory modalities enhances brain-machine interface control. The Journal of Neuroscience 30:16777–16787.

Sun W, Dan Y (2009) Layer-specific network oscillation and spatiotemporal receptive field in the visual cortex. Proceedings of the National Academy of Sciences of the United States of America 106:17986–17991.

Sussillo D, Churchland MM, Kaufman MT, Krishna V (2015) A neural network that finds naturalistic solutions for the production of muscle activity. Nature Neuroscience 18:1025–1033.

Sussillo D, Nuyujukian P, Fan JM, Kao JC, Stavisky SD, Ryu SI, Shenoy K V (2012) A recurrent neural network for closed-loop intracortical brain-machine interface decoders. Journal of Neural Engineering 9:26027.

Takahashi K, Saleh M, Penn RD, Hatsopoulos NG (2011) Propagating Waves in Human Motor Cortex. Frontiers in Human Neuroscience 5:40.

Tan H, Jenkinson N, Brown P (2014) Dynamic neural correlates of motor error monitoring and adaptation during trial-to-trial learning. The Journal of Neuroscience 34:5678–5688.

Tan H, Wade C, Brown P (2016) Post-Movement Beta Activity in Sensorimotor Cortex Indexes Confidence in the Estimations from Internal Models. The Journal of Neuroscience 36:1516–1528.

Taylor DM, Tillery SIH, Schwartz AB (2002) Direct cortical control of 3D neuroprosthetic devices. Science 296:1829–1832.

Tennant KA, Adkins DL, Donlan NA, Asay AL, Thomas N, Kleim JA, Jones TA (2011) The organization of the forelimb representation of the C57BL/6 mouse motor cortex as defined by intracortical microstimulation and cytoarchitecture. Cerebral Cortex 21:865–876.

Thompson KG, Biscoe KL, Sato TR (2005) Neuronal Basis of Covert Spatial Attention in the Frontal Eye Field. The Journal of Neuroscience 25:9479–9487.

Toppi J, Risetti M, Quitadamo LR, Petti M, Bianchi L, Salinari S, Babiloni F, Cincotti F,

Mattia D, Astolfi L (2014) Investigating the effects of a sensorimotor rhythm-based BCI training on the cortical activity elicited by mental imagery. Journal of Neural Engineering 11:35010.

Torene S, Brincat SL, Salazar-Gomez AF, Jia N, Panko M, Saligrama V, Miller EK, Guenther FH (2013) Adaptive decoding of eye movements with a simple recurrent artificial neural network. In: 2013 Neuroscience Meeting Planner, pp.373.09.

Torene S, Ritt JT, Guenther FH (2015) Brain machine interface control through neurofeedback guided beta rhythm modulation. In: Neuroscience Meeting Planner, pp.713.04.

Torrecillos F, Alayrangues J, Kilavik BE, Malfait N (2015) Distinct modulations in sensorimotor postmovement and foreperiod β-band activities related to error salience processing and sensorimotor adaptation. The Journal of Neuroscience 35:12753–12765.

Tremblay S, Pieper F, Sachs A, Martinez-Trujillo J (2015) Attentional Filtering of Visual Information by Neuronal Ensembles in the Primate Lateral Prefrontal Cortex. Neuron 85:202–215.

Tzagarakis C, Ince NF, Leuthold AC, Pellizzer G (2010) Beta-band activity during motor planning reflects response uncertainty. The Journal of Neuroscience 30:11270–11277.

Velliste M, Kennedy SD, Schwartz AB, Whitford AS, Sohn J-W, McMorland AJC (2014) Motor Cortical Correlates of Arm Resting in the Context of a Reaching Task and Implications for Prosthetic Control. The Journal of Neuroscience 34:6011–6022.

Velliste M, Perel S, Spalding MC, Whitford AS, Schwartz AB (2008) Cortical control of a prosthetic arm for self-feeding. Nature 453:1098–1101.

Vidaurre C, Blankertz B (2010) Towards a cure for BCI illiteracy. Brain Topography 23:194–198.

Vidaurre C, Sannelli C, Müller K-R, Blankertz B (2011) Machine-learning-based coadaptive calibration for brain-computer interfaces. Neural Computation 23:791–816.

Vidaurre C, Schloogl A, Cabeza R, Scherer R, Pfurtscheller G (2006) A fully on-line adaptive BCI. IEEE Transactions on Bio-Medical Engineering 53:1214–1219.

Voigts J, Sakmann B, Celikel T (2008) Unsupervised whisker tracking in unrestrained behaving animals. Journal of Neurophysiology 100:504–515.

Wahnoun R, He J, Helms Tillery SI (2006) Selection and parameterization of cortical neurons for neuroprosthetic control. Journal of Neural Engineering 3:162–171.

Wander JD, Blakely T, Miller KJ, Weaver KE, Johnson LA, Olson JD, Fetz EE, Rao RPN, Ojemann JG (2013) Distributed cortical adaptation during learning of a brain – computer interface task. Proceedings of the National Academy of Sciences of the United States of America 110:10818–10823.

Wang D, Zhang Q, Li Y, Wang Y, Zhu J, Zhang S, Zheng X (2014) Long-term decoding stability of local field potentials from silicon arrays in primate motor cortex during a 2D center out task. Journal of Neural Engineering 11:36009.

Wang Z, Gunduz A, Brunner P, Ritaccio AL, Ji Q, Schalk G (2012) Decoding onset and direction of movements using Electrocorticographic (ECoG) signals in humans. Frontiers in Neuroengineering 5:15.

Wessberg J, Stambaugh CR, Kralik JD, Beck PD, Laubach M, Chapin JK, Kim J, Biggs SJ, Srinivasan MA, Nicolelis MAL (2000) Real-time prediction of hand trajectory by ensembles of cortical neurons in primates. Nature 408:361–365.

Whittingstall K, Logothetis NK (2009) Frequency-band coupling in surface EEG reflects spiking activity in monkey visual cortex. Neuron 64:281–289.

Williams JJ, Rouse AG, Thongpang S, Williams JC, Moran DW (2013) Differentiating closed-loop cortical intention from rest: building an asynchronous electrocorticographic BCI. Journal of Neural Engineering 10:46001.

Witham CL, Wang M, Baker SN (2007) Cells in somatosensory areas show synchrony with beta oscillations in monkey motor cortex. The European Journal of Neuroscience 26:2677–2686.

Witte M, Kober SE, Ninaus M, Neuper C, Wood G (2013) Control beliefs can predict the ability to up-regulate sensorimotor rhythm during neurofeedback training. Frontiers in Human Neuroscience 7:478.

Wolpaw JR, McFarland DJ (2004) Control of a two-dimensional movement signal by a noninvasive brain-computer interface in humans. Proceedings of the National Academy of Sciences of the United States of America 101:17849–17854.

Wood F, Black MJ, Vargas-Irwin C, Fellows M, Donoghue JP (2004) On the variability of manual spike sorting. IEEE Transactions on Biomedical Engineering 51:912–918.

Wood G, Kober SE, Witte M, Neuper C (2014) On the need to better specify the concept of “control” • in brain-computer-interfaces/neurofeedback research. Frontiers in Systems Neuroscience 8:171.

Wu W, Hatsopoulos NG (2008) Real-time decoding of nonstationary neural activity in motor cortex. IEEE Transactions on Neural Systems and Rehabilitation Engineering 16:213–222.

Wyrwicka W, Sterman MB (1968) Instrumental conditioning of sensorimotor cortex EEG spindles in the waking cat. Physiology & Behavior 3:703–707.

Xu N, Harnett MT, Williams SR, Huber D, O'Connor DH, Svoboda K, Magee JC (2012) Nonlinear dendritic integration of sensory and motor input during an active sensing task. Nature 492:247–251.

Zaepffel M, Trachel R, Kilavik BE, Brochier T (2013) Modulations of EEG beta power during planning and execution of grasping movements. PLoS ONE 8:e60060.

Zelaznik HN, Hawkins B, Kisselburgh L (1983) Rapid visual feedback processing in single-aiming movements. Journal of Motor Behavior 15:217–236.

CURRICULUM VITAE

Education

Current	**Ph.D. Candidate**, Boston University, Computational Neuroscience
2004	**B.S.**, University of Maryland, Computer Science

Research Experience

2010-Present	**Ph.D. Candidate** Advisors: Dr. Frank H. Guenther and Dr. Jason T. Ritt Graduate Program for Neuroscience Boston University, Boston MA
2002-2004	**Research Internship** Advisor: Dr. Bruce Weber U.S. Army Research Laboratory, Adelphi MD

Professional Positions

2010	**Professional Services Engineer** Courion Corporation, Westborough MA
2006-2010	**Systems Engineer** **Systems Administrator II** **Systems Administrator I** Web.com, Atlanta GA
2004-2006	**Jr. Messaging Administrator** Public Broadcasting Service, Alexandria VA

Teaching Experience

2011	**Teaching Fellow**, NE203, Principles of Neuroscience, Boston University

Fellowships

2014-2015	Center of Excellence for Learning in Education, Science, and Technology, National Science Foundation
2012-2014	Computational Neuroscience Fellowship, National Institutes of Health
2010-2012	Traineeship in Neuroscience, National Institutes of Health

Awards and Honors

2013	Neuroscience 2013 “Hot Topic”, San Diego CA

Academic Service

2014-2015 **Neural Prosthesis Lab Journal Club Czar,** Boston University
2014 **Mentor**, Incoming GPN Student, Boston University
2014 **Presenter**, Regular Expressions at Grad Hackathon, Boston MA
2013-2015 **Moderator**, GPN Retreat Breakout Discussion Groups, Essex MA
2013, 2015 **Panelist**, CELEST Career Panel, Boston University
2012-2015 **Election Organizer**, Neuroscience Graduate Student Organization Elections, Boston University
2012-2014 **Organizer**, Five Minute Talks, Boston University
2012-2013 **Vice President**, Computational Neuroscience Student Organization, Boston University
2012 **Presenter**, Superparamagnetic Clustering, Boston University
2011, 2013-16 **Presenter**, U-Design Demonstration, Boston University

Peer-Reviewed Publications

2010 Iskow RC, McCabe MT, Mills RE, Torene SB, Pittard WS, Neuwald AF, Van Meir EG, Vertino PM, and Devine SE. Natural mutagenesis of human genomes by endogenous retrotransposons. *Cell.* Jun 25;141(7):1253-61.

Posters and Talks

2015 Torene S, Ritt JT, Guenther FH. Brain machine interface control through neurofeedback guided beta rhythm modulation. Neuroscience. Society for Neuroscience, Chicago IL. 2015. Poster.

2015 Torene S, Ritt JT, Guenther FH. Brain machine interface control through neurofeedback guided beta rhythm modulation. 8th Annual inter-Science of Learning Center Student and Post-Doc Conference. University of California, San Diego, La Jolla, CA. 2015. Poster.

2015 Torene S, Ritt JT, Guenther FH. Brain machine interface control through neurofeedback guided beta rhythm modulation. NeuroHAM. Boston University, Boston MA. 2015. Poster.

2015 Torene S, Ritt JT, Guenther FH. Brain machine interface control through neurofeedback guided beta rhythm modulation. Graduate Program for Neuroscience Recruitment Reception. Boston University, Boston MA. 2015. Poster.

2015 Torene S, Ritt JT, Guenther FH. Brain machine interface control through neurofeedback guided beta rhythm modulation. Cognitive Rhythms Collaborative Retreat. Boston University, Boston MA. 2015. Poster.

2014 Torene S, Brincat SL, Panko M, Jia N, Salazar-Gomez AF, Saligrama V, Miller EK, Guenther F. Adaptive decoding of eye movements with a simple recurrent artificial neural network. 20th Annual Henry I. Russek Student Achievement Day. Boston University, Boston MA. 2014. Poster.

2013 Adaptive Decoding of Eye Movements with a Simple Recurrent Artificial Neural Network. Graduate Program for Neuroscience Retreat. Essex Conference Center, Essex MA. 2013. Invited talk.

2013 Torene S, Brincat SL, Panko M, Jia N, Salazar-Gomez AF, Saligrama V, Miller EK, Guenther F. Adaptive decoding of eye movements with a simple recurrent artificial neural network. Neuroscience. Society for Neuroscience, San Diego CA. 2013. Poster.

2013 Torene S, Brincat SL, Panko M, Jia N, Salazar-Gomez AF, Saligrama V, Miller EK, Ritt JT, Guenther F. Adaptive decoding of eye movements with a simple recurrent artificial neural network. 17th International Conference on Cognitive and Neural Systems. Boston University, Boston MA. 2013. Poster.

2013 Torene S, Brincat SL, Panko M, Jia N, Salazar-Gomez AF, Saligrama V, Miller EK, Ritt JT, Guenther F. Adaptive decoding of eye movements with a simple recurrent artificial neural network. Graduate Program for Neuroscience Recruitment Reception. Boston University, Boston MA. 2013. Poster.